THE NATIONAL MARITIME DOMAIN AWARENESS ARCHITECTURE PLAN

A NATIONAL MARITIME INFORMATION SHARING ENVIRONMENT (MISE) IMPLEMENTED THROUGH COMMON DATA STANDARDS AND ARCHITECTURAL UNDERSTANDING

February 2016

THIS PAGE INTENTIONALLY LEFT BLANK

ABOUT THIS DOCUMENT

The National Maritime Domain Awareness Architecture Plan describes a process for sharing maritime information. To provide the appropriate level of detail and technical specificity to the full spectrum of readers, the document is divided into multiple sections. They include:

Executive Summary

The *Executive Summary* is for senior leaders who want an overview of the plan. It summarizes major concepts and business processes.

Concept

The *Concept* is for mid-level managers and senior technical managers who need to understand, with some detail, the objectives, features and business processes of the plan.

Framework

The *Framework* is for technical managers or technology implementers who need to understand, with greater detail, the technical implementation and specifications of the plan.

Appendices

The *Appendices* are a collection of architectural and technical documents which are relevant to, but not otherwise included in, the National Maritime Domain Awareness Architecture Plan. Each appendix represents the state of the technology, process or definition at the time this plan was published.

Attachments

The *Attachments* are a collection of documents that describe the National Information Exchange Model – Maritime (NIEM-M), and its implementation within the maritime Community, at the time this plan was published. While NIEM-M is not part of the plan, the definition of common data standards is integral to the success and implementation of this plan.

This plan does not alter existing constitutional or statutory authorities or responsibilities of department and agency heads to carry out operational activities or to exchange information. The Maritime Security Interagency Policy Committee (MSIPC), its successor, or its designated representative, will review the National Maritime Domain Awareness Architecture Plan upon significant changes to PPD-18 or the NSMS or every five years, whichever occurs first. The Maritime Domain Awareness (MDA) Executive Steering Committee (ESC) may modify the NMDAAP appendices as required.

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

A NATIONAL MARITIME INFORMATION SHARING ENVIRONMENT (MISE) IMPLEMENTED THROUGH COMMON DATA STANDARDS AND ARCHITECTURAL UNDERSTANDING

"It is a national priority to efficiently, effectively, and appropriately share and safeguard information so any authorized individual (Federal, state, local, tribal, territorial, private sector or foreign partner) can prevent harm to the American people and protect national security. The Strategy points toward a future in which information supports national security decision-making by providing the right information, at any time, to any authorized user, restricted only by law or policy, not technology; and where safeguarding measures, to include a comprehensive regimen of accountability, prevent the misuse of the information."[1]

This passage from the *National Strategy for Information Sharing and Safeguarding* (the National Strategy) establishes the National vision for information sharing. The National Strategy goes on to identify three core information sharing principles:

1. Information as a National Asset
2. Information Sharing and Safeguarding Requires Shared Risk Management
3. Information Informs Decision Making

Guided by these principles and as directed by the National Maritime Domain Awareness Plan (NMDAP), the Department of Defense Executive Agent for Maritime Domain Awareness (DoD EA for MDA) led the effort to, *"establish a national-level MDA enterprise architecture to include: Enterprise, web-centric, cloud-based, information and services; Common data standards; and, Data access policy."*[2, 3] The resulting Maritime Information Sharing Environment (MISE) is captured in this plan. The MISE provides an internet accessible, unclassified information sharing capability where data providers and consumers manage and share maritime information through common data definitions and security attributes.

Being mindful of the principles, and priority objectives, identified in the National Strategy, the objectives in the NMDAP, and with full understanding that each participating agency has its own operational constraints; the MISE defines an operational and technical framework that enables information sharing by leveraging existing programs and systems. Within the U.S., agencies and organizations with maritime interests have their own requirements, authorities, information infrastructures, and resources. The MISE acknowledges each of those constraints and defines a service oriented architectural approach that allows participation while protecting individual information and resources. Data and analytical products are shared via common data

[1] *National Strategy for Information Sharing and Safeguarding*, December 2012, Vision Statement
[2] *National Maritime Domain Awareness Plan*, December 2013, Implementation Plan
[3] This also aligns with the National Strategy's Priority Objectives to include 1, 3, 4, and 11.

standards with access controls that allow data providers to manage their information within the constraints of their respective authorities or regulations.

The MISE employs industry standard service protocols that, in most cases, are already supported by participating agencies. The MISE *does not* define how agencies execute their missions; manage their infrastructure, or the products they develop. The MISE *does* define how the results of those mission activities or information products are made available to others. Through the definition of data standards within the National Information Exchange Model-Maritime (NIEM-M)[4], the MISE provides a common vocabulary in four initial focus areas: Vessel Positions, Advance Notice of Arrival, Indicators and Notifications, and Maritime Operational Threat Response (MOTR)[5]. While only four focus areas are defined in the initial effort, the standards and processes defined by the MISE are designed to be reusable and extensible to support future information sharing products and partners.

Currently, the concepts, capabilities, and standards outlined in this plan are being implemented across the U.S. Government. The Office of Naval Intelligence has adopted NIEM-M as a standard for future maritime data services. The Navy Data Engineering Sciences Center (DESC) is a DoD organization supporting NIEM-M development. The technologies described in this plan have also been tested and are currently implemented and sharing maritime data within multiple programs. The Coast Guard Interagency Operations Center project, the Air Force Maritime Domain Awareness System at MacDill AFB in Tampa, Florida and U.S. SOUTHCOM's Caribbean Sensor Information Integration are three programs operating today.

For the MISE to fulfill the principles outlined in the National Strategy, additional steps are required. The continued development of standards and the implementation of services from authoritative data sources are important first steps. Critical to the success is developing confidence in the processes and trust in the partnerships. The DESC has implemented the MISE in a commercial cloud environment for use by international, interagency or industry partners to develop and test the web services required to build that confidence and trust.

Information sharing is a mandated activity that directly affects National Security. Information sharing efforts must be conducted consistent with legal and policy constraints. The MISE defines a low cost, implementable solution while providing mechanisms to accommodate legal and policy constraints. While the MISE may not support the information sharing requirements of every agency, it will support the processing and management of the majority of the unclassified maritime information the nation collects, processes, and shares to support and improve our National Security in the maritime domain.

[4] Data Standards and NIEM also align to the National Strategy's Priority Objectives 3 and 10.

[5] While MOTR was designated as an initial focus area, the MOTR process was not used in the Architecture Plan.

TABLE OF CONTENTS

THIS PAGE INTENTIONALLY LEFT BLANK

CONCEPT

1. Background

The *National Strategy for Information Sharing and Safeguarding* (the National Strategy) and the *National Maritime Domain Awareness Plan (NMDAP)* describe an information sharing architecture founded upon data-centric principles to provide a secure, collaborative, information-sharing environment (ISE). The described ISE recognizes information as a national asset and safeguarding information is a shared responsibility.

In October 2010, the National Maritime Domain Awareness (MDA) governing body, the MDA Executive Steering Committee[6], assigned the development of the National Maritime Domain Awareness Architecture Plan.

The goal of this document is to define the features, processes and technologies required to implement the Maritime Information Sharing Environment (MISE). The MISE is the body of the National Maritime Domain Awareness Architecture Plan (also referred to as the Architecture Plan) and is the term that embodies the participants and functions described herein.

2. Overview

The Architecture Plan describes how maritime information providers may share and protect their information within the MISE. The Global Maritime Community of Interest (GMCOI) is broad and includes Federal, state, local, tribal, territorial and international partners. Maritime information must be integrated and shared across the GMCOI to gain an effective understanding of the maritime domain. The Architecture Plan is a solution to mitigate and manage gaps between the different information sharing systems and user requirements.

Being mindful of the principles identified in the National Strategy, the objectives in the NMDAP and with full understanding that each participating agency has its own operational constraints; the MISE defines an operational and technical framework that enables information sharing by leveraging existing programs and systems to the greatest extent possible. Agencies and organizations with maritime interests have their own requirements, authorities, information infrastructures, and resources. The MISE accounts for each of those constraints and defines a service oriented architectural approach that allows participation while protecting individual information and resources. Data and analytical products are shared via common data standards with access controls that enable data providers to manage sharing as defined by their respective authorities or regulations.

[6] The MDA Executive Steering Committee members are Senior Executive Service or Flag Officers from the Departments of Defense, Homeland Security and Transportation and the Office of the Director of National Intelligence.

To meet the goals and objectives described in this plan and outlined in the National Strategy and NMDAP, the MISE is comprised of four functional parts. They include:

- Trusted systems and their users
- NIEM-Maritime data standards
- Common attributes for access control
- Information Sharing Infrastructure (ISI)

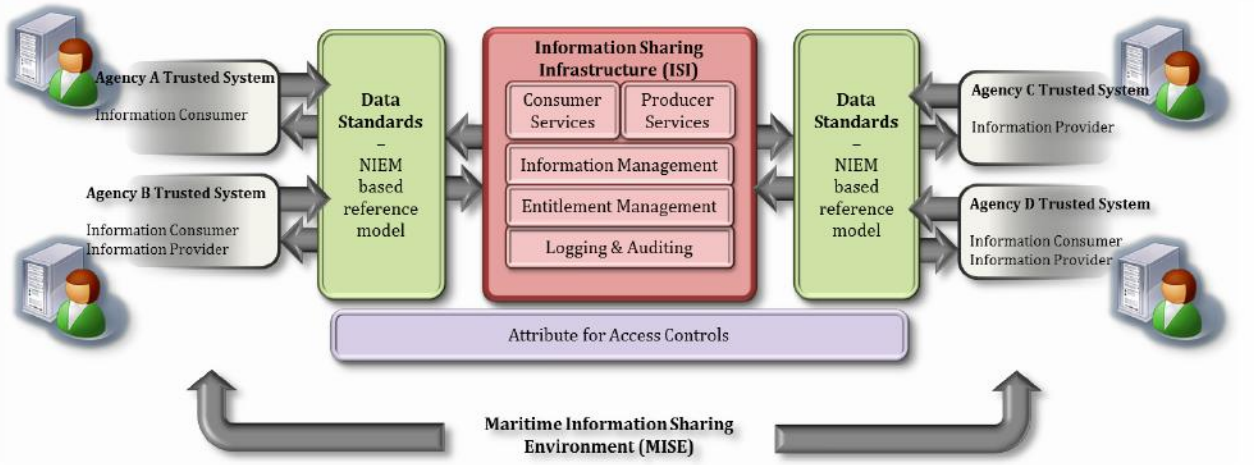

Figure 1: Operational view of the MISE

The breadth of all possible maritime information sharing is greater than can be described in a single document or undertaken in a single effort. To scope the initial effort and manage the volume of information contained in this plan, four focus areas were identified for the initial effort. They include:

1. Notice of Arrival (NOA)[7]. A 96-hour advance notice that vessels over 300 gross tons and inbound to US ports are required to submit which lists vessel, crew, passenger, and cargo information.
2. Indicators and Notifications (IAN). Indicators are information used to inform or contribute to an analytical process. Notifications include warnings of a possible event and alerts about the execution of an event.
3. Positions (POS). A geospatial position, course, heading, speed, and status of a vessel at a given time. A series of position reports can be combined to produce track information.
4. Maritime Operational Threat Response (MOTR)[8]. A cross-departmental conference convened as necessary to coordinate response to a maritime threat.

To ensure proper data security and entitlement management, the Architecture Plan employs an attributes-based sharing policy that defines an Information Access Policy (IAP) process. The

[7] Notice of Arrival (NOA) is interchangeable with Advanced Notice of Arrival (ANOA) per IEPD.
[8] While MOTR was designated as an initial focus area, the MOTR process was not used in the Architecture Plan.

IAP is a predefined set of rules, represented using common attributes, that governs which users can access what information. The attributes-based sharing policy provides a level of assurance to information providers that information is properly handled and access is only granted to authorized users.

3. Understanding Information Sharing

Information sharing is *not* about moving information between data providers and data consumers. It's *not* about defining common standards or security protocols. It's *not* about sharing for the sake of sharing. Information sharing *is* about improving our business processes to execute our defined missions to the best extent possible. To improve our business processes and better execute our missions, we need to gather as much relevant information as possible or provide information products that are the result of our analytical missions.

The Architecture Plan describes a repeatable information sharing process governed by common rules for information management. The Architecture Plan uses a common vocabulary to describe the information that is shared and why we're sharing it. "Why" is the business, mission or operational reason we expend resources to implement the sharing services. The "why" may be the most important but most overlooked part of any information sharing effort.

Finally, information sharing is about more than gaining access to large amounts of unprocessed or raw data. It's about defining and sharing the products of data processing, analysis or fusion. The Architecture Plan amplifies the idea that information sharing is not about discrete data elements or sharing raw sensor based data. Information sharing must move beyond sharing individual data elements and include the idea of *data products* that are the result of analytical processes or data processing capabilities and provide operational relevance. While a contact from a radar or acoustical sensor may be of value, that same contact correlated or fused to other data providing vessel name and characteristics is more valuable to the community.

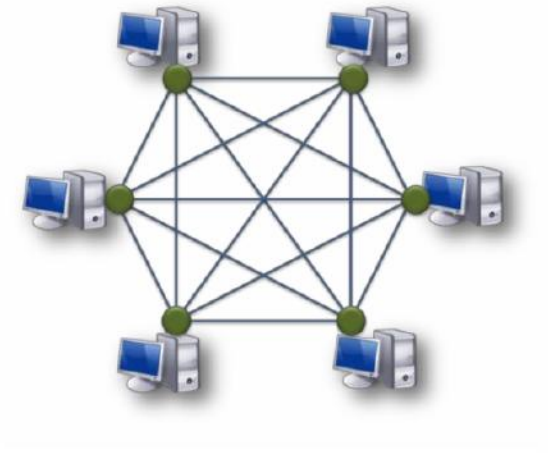

Figure 2: Point-to-point sharing

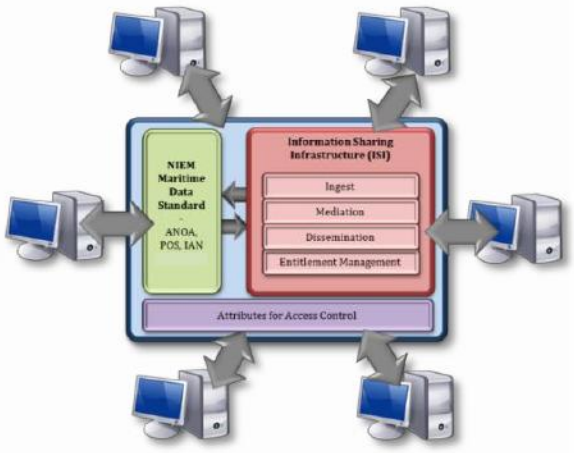

Figure 3: MISE services based sharing

Figure 2 represents the complexity of current point-to-point sharing activities. Each participating system must establish a discrete connection point to every other system. This type of sharing environment requires every participant to update its connection point whenever a new partner joins the environment. Figure 3 represents MISE services based sharing. Each

participating system maintains only one connection point to the MISE. When a new participant joins the MISE, no changes are required by existing members. The ISI continues entitlement management for the new and existing information according to the common security attributes.

3.1. INFORMATION SHARING BUSINESS PROCESS

Successful information sharing activities are the result of operational, information and technological understanding achieved through a well-defined and routinely implemented process. The Architecture Plan describes a multistep information sharing process.

1. Describe the operational use case being supported by the information sharing.
2. Identify the specific data elements required to support the use case.
3. Develop a standard definition, model or product for the information to be shared.
4. Identify any legal or policy driven constraints on the information.
5. Implement appropriate controls to ensure proper entitlement management.
6. Implement and monitor the sharing service.

Each of the six steps will be discussed in greater detail throughout the Architecture Plan.

3.2. NATIONAL DATA PRODUCTS: A NATIONAL ASSET

The first of the core principles outlined in the National Strategy is *Information as a National Asset*. While information must be safeguarded, it first must be made available. Each data provider has legal, contractual or policy reasons that may prevent some of the data they maintain or process from being shared. However; the *products* of their analytical and management process may have a broader release. It's the obligation of those with data products to make those products discoverable and sharable; even if that means the release of multiple versions with different security controls. Leveraging the security markings, such as the attributes in NIEM-M based exchanges; a single data model can be used to support multiple versions of a product. The idea of *Information as a National Asset* can be realized through national information products. Information products, derived from national maritime programs like the Department of Homeland Security Maritime Information Global Network (MAGNet), make the results of complex analysis or data fusion processes available to consumers who might have less mature analytical or IT processes. The availability of information products from these types of programs relieve other Agencies or programs from duplicating the analysis thus providing cost saving and resulting in a common maritime understanding. Products like *Consolidated Vessel Information & Security Reporting*[9] can be provided via a simple Google Earth interface yet give the consumer a deep understanding of the maritime domain.

[9] *Consolidated Vessel Information & Security Reporting* (CVISR), as defined by the Department of Defense, is a standard definition of maritime understanding to support national security. Details about CVISR and the CVISR Information Exchange Package Documentation (IEPD) can be found in the CVISR Exchange Summary Attachment.

3.3. USE CASES

The first step to any information sharing effort is the definition of the operational context; the reason why the sharing action should even take place. Broad ideas like "improve national security" or "stop illegal activities" are reasons that have been used to define the "why" of information sharing but rarely provide the fidelity required to define the specific data elements that must be shared. These broad descriptions lead to long and often unproductive discussions about what specific data should be shared, who has the authority to provide or consume the information, and what it will be used for once it is shared. To ensure both data providers and consumers are able to work within their own legislative or policy bounds, the "why" of the sharing must be as specific as the data elements that are shared.

Use cases have been defined for each of the initial focus areas within the Architecture Plan. The use cases were used to identify the specific data elements contained in the supporting model and how the information would be shared. In lieu of listing the entire set of use cases, they are represented as a graphic (Figure 4), depicting common blocks and data elements.

3.4. DATA STANDARDS

Within the Architecture Plan, "data standards" are the common vocabulary used to describe the shared data objects and the models themselves. The National Information Exchange Model[10] (NIEM) and NIEM-Maritime serve as the reference library used to develop the specific data objects to be shared. The information contained in NIEM is stored in an eXtensible Markup Language (.xml) format. The technical representation of .xml is complex and it can be difficult to track the relationship of NIEM and NIEM-Maritime components. To facilitate discussion and demonstrations with business and process managers, a graphical representation of each data model has been developed.

Figure 4 is a graphical representation of the logical data model for Position data. Each Position record is defined with *Movement*, *Position*, and *Vessel Identification* components as defined in NIEM. *Record Metadata* provides security and record control information.

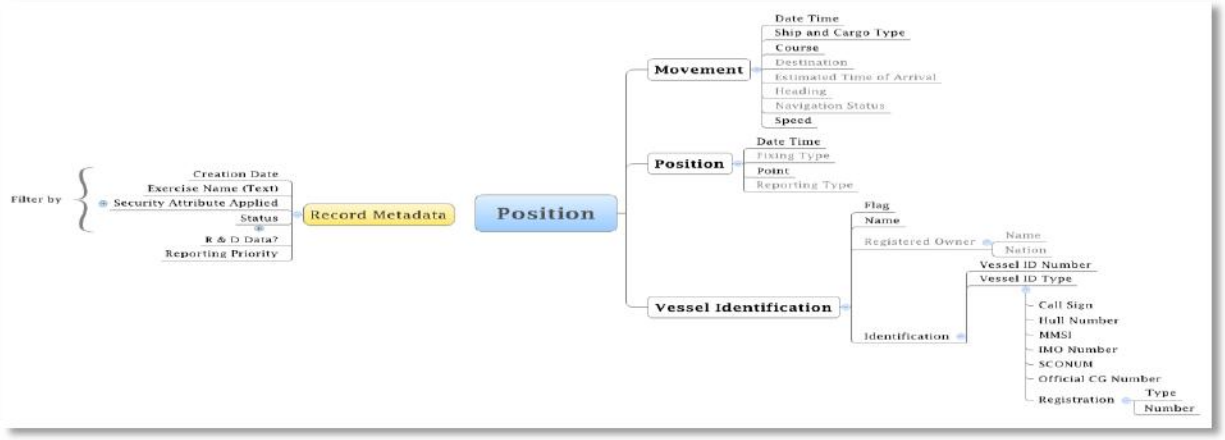

Figure 4: Position logical data model

[10] NIEM, a Department of Justice and Department of Homeland Security led effort, is the National guideline for sharing data between United States government agencies and their partners.

Figure 5 is a graphical representation of the request use case for position data. Positions can be requested by *Geo Area* or *Vessel*, both with specific selection and bounding variables, and filtered by *Record Metadata*. Examples of request strings are available in the *Implementation Guide* which can be found in Appendix A.

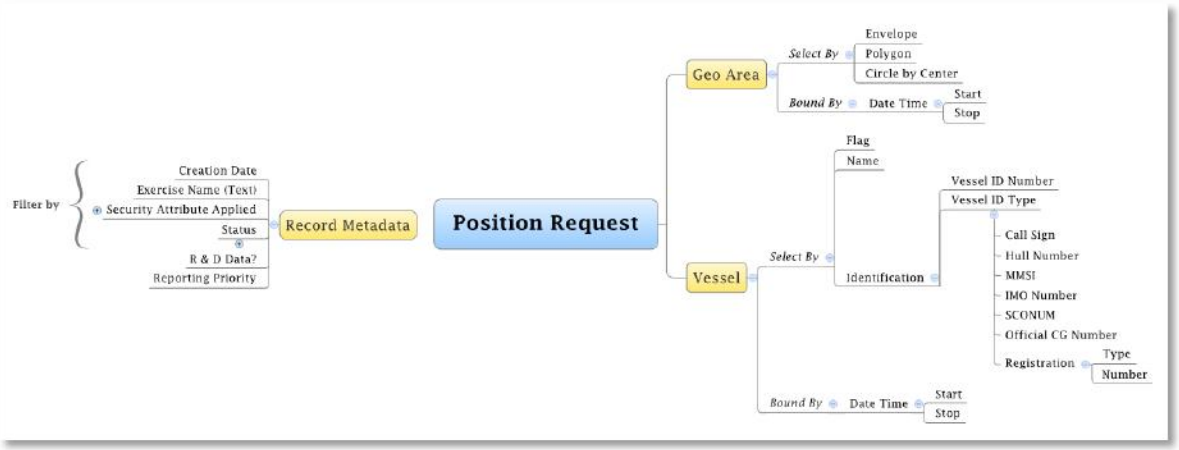

Figure 5: Position Request Use Cases

3.5. SECURITY ATTRIBUTES

Security attributes and record metadata are the elements used to provide access control for the data object. Security attributes are applied at the message, record and user level.

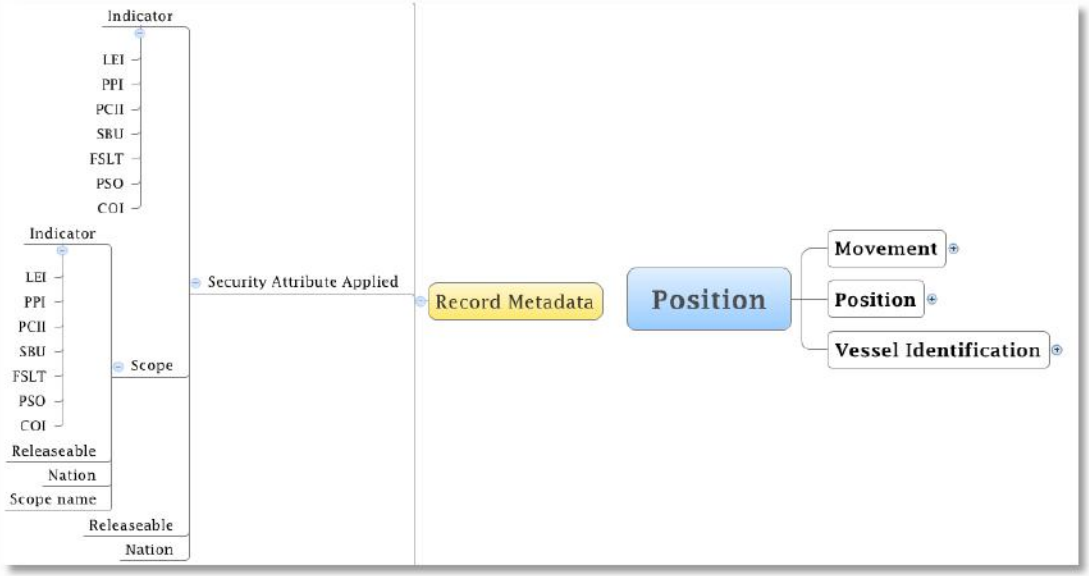

Figure 6: Record level security attributes

As shown in figure 6, security attributes include *Indicator, Releasable, Nation, and Scope*. Currently, indicators are Law Enforcement (LEI), Protected (PPI), Protected Critical Infrastructure Information (PCII), Sensitive but Unclassified (SBU), Federal State Local and Tribal (FSLT), Private Sector Only (PSO), and Community of Interest (COI), with COI representing all the

participants of the MISE. *Releasable* indicates if the record is publicly releasable and *Nation* is a list of nations that can receive the record. *Scope*, which is made up of a scope name and all the previous listed attributes is designed to allow auxiliary control of the record. Security attributes are an integral component within the NIEM-Maritime data model but pose a separate set of parameters, specifically outlining the level of sensitivity for a participating agency's control measures. Full details about *Attribute Specifications* and their implementation can be found in Appendix B.

4. The MISE

The Maritime Information Sharing Environment, (MISE), is the regulated application of the Architecture Plan. The term MISE expresses the type, and functionality of the information participants desire to share, directly in reflection to the guidelines within the Architecture Plan. MISE consists of:

- Trusted systems and their users
- Information Sharing Infrastructure
- Data standards NIEM based reference model
- Security attributes for entitlement management

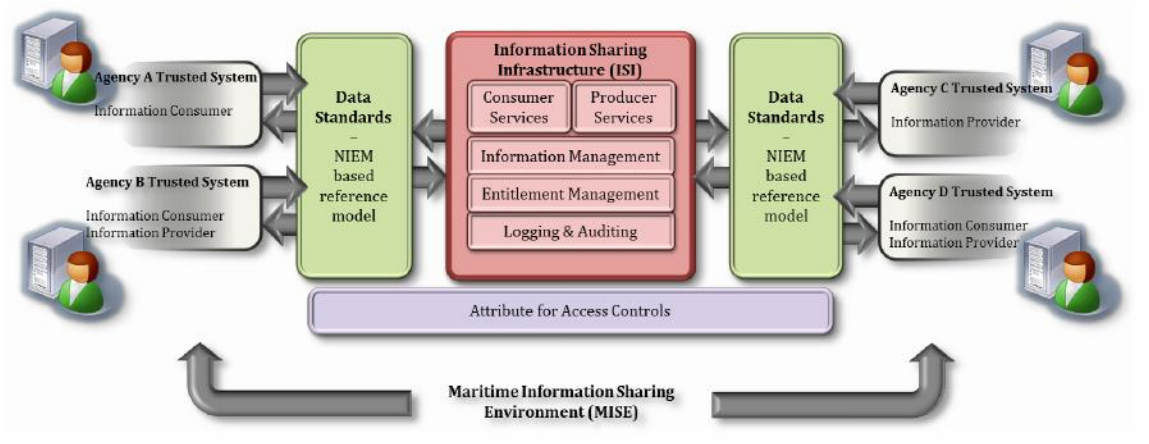

Figure 7: Maritime Information Sharing Environment

4.1. INFORMATION SHARING INFRASTRUCTURE (ISI) FEATURES

The Information Sharing Infrastructure (ISI) is the central hub for information and its distribution within the MISE; it is the component that provides the information sharing service to the trusted systems. Participating agencies will be associated to the MISE via their trusted system which will allow them to share or retrieve information by connecting and interacting with the ISI. Trusted systems can serve as information providers, consumers, or both. Trusted systems cannot and do not interact directly with, or access, other trusted systems; rather they interact only through the ISI or central exchange hub which processes and distributes each

request individually. A trusted system is a participant's new or legacy IT system that allows a user to process, manage or view available maritime information for publication to or exchange with the ISI. A full description of trusted system requirements is outlined in Appendix F.

Information is captured and processed through the ISI into three functional categories: Entitlement Management, Information Management, and Logging/Auditing. The figure below depicts the information flow within the ISI, and how individual features are broken out across each functional area.

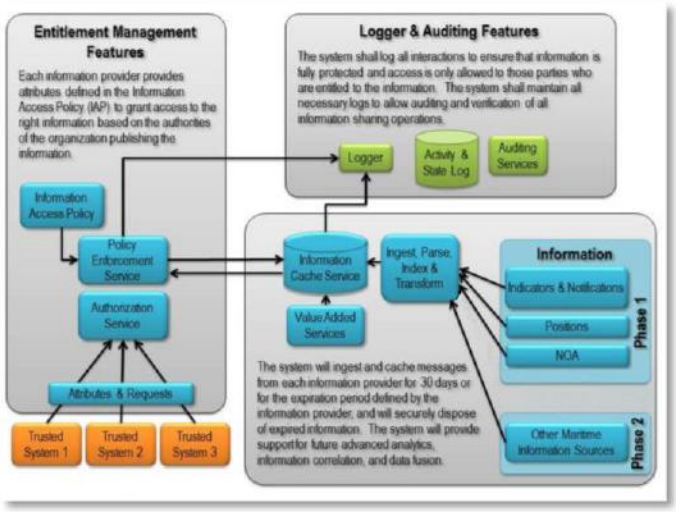

Figure 8: Information Sharing Infrastructure Features

A full description of the ISI is contained throughout the appendices.

4.2. MISE FRAMEWORK

The Architecture Framework, described in detail in the *Framework* section is comprised of four architectural views: Data, Services, Security, and Technical Operations as depicted in the following figure.

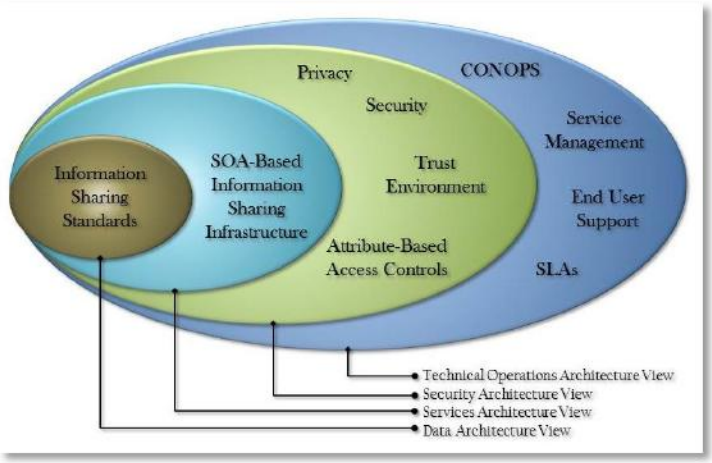

Figure 9: The Four Architectural Views of the MISE Framework

Data Architecture View

The Data Architecture View is the center of the architecture framework. It includes the standard characteristics and vocabulary used to define standards for information sharing. These standards are a part of the NIEM.

Services Architecture View

The Services Architecture View builds upon the Data Architecture View, and is the foundation for the MISE. It includes the information sharing infrastructure that provides common service interfaces for trusted systems to share information.

Security Architecture View

The Security Architecture View ensures that the shared information is protected. It includes the common security attributes and the association of the trusted systems.

Technical Operations View

The Technical Operations View defines how the framework as a whole is managed and operated. It includes governance and agreements used to manage the MISE.

Below is an example of how MISE services are used together across all four architectural views of the framework to deliver value and insight to users.

In a given region surrounding a port, a graphical display shows incoming and outgoing vessels. Users of the system can click on a vessel and retrieve detailed information about the vessel. Users also have the ability to display tracks showing the path taken by the vessel to reach its current position. The notional graphic below shows how a simple interface can be used to receive large amounts of data from the MISE. Representative Level of Awareness information is shown.

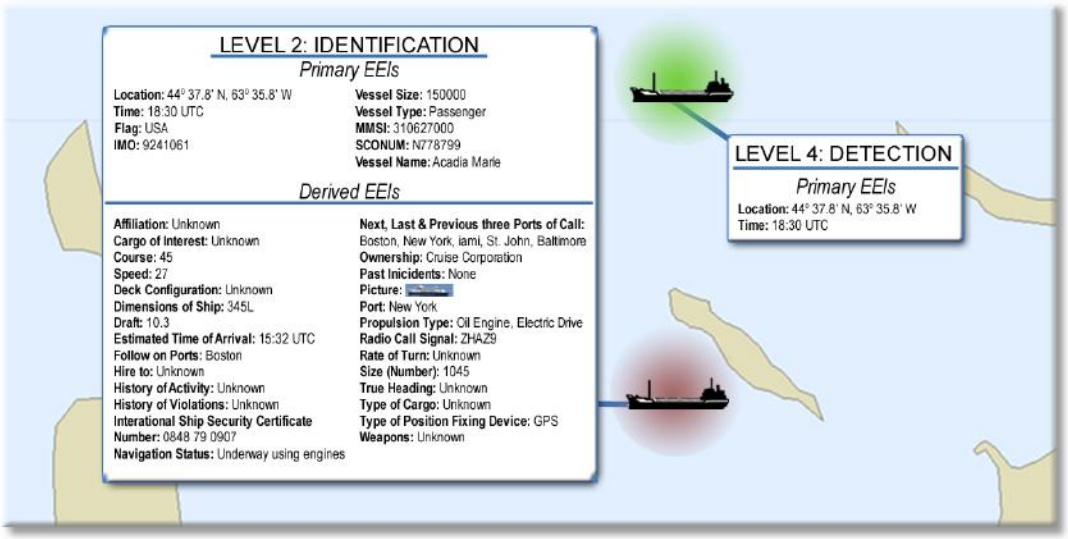

Figure 10: Representational of MISE Services

4.3. MISE EVOLUTION

The Architecture Plan, including the implementation of the ISI and definition of data standards and processes, is partially bounded by the technical and process maturity of the participating maritime community. As the community and its processes mature, the trust of data providers

and their products will grow. Improvements in the level of trust between the MISE participants and technical capabilities will encourage more sharing and additional services. With time, the MISE will provide greater services and value added products. As that happens, greater fidelity in the security attributes will be required. The MISE has been designed and developed to deliver the needs of today's user community while allowing the flexibility to grow with the maturity of that community.

5. Security

To many participants of the MISE, security is the most important aspect of implementation. For the Architecture Plan, security is designed to be implemented the same across the MISE yet controlled by individual data providers.

Security is based upon a common set of attributes. The three categories for attributes are:

1. Entity Attributes: Attributes associated to a trusted system.
2. User Attributes: Attributes associated to a specific user.
3. Data Attributes: Attributes associated to a data record.

When a data provider publishes information, they will provide the appropriate security attributes to the ISI for each record. When a trusted system requests information on behalf of a user, the ISI will ensure the attributes of the user match those of the data provider and only share the appropriate marked information. Since the attributes are asserted at the time the information is published, two direct benefits are achieved:

1. The attributes are immediately applied to the data. No additional system changes are required.
2. The attributes are associated to the information and passed to data consumers. This ensures the data consumer understands the protection and management requirements of the data provider.

5.1. SECURITY OVERVIEW

In order for users to feel comfortable that their information is protected, the Architecture Plan defines the MISE trust fabric. The MISE trust fabric is an .xml document that is digitally signed by the MISE Certificate Authority (CA). The trust fabric is a critical document that describes each trusted system in Security Assertion Markup Language (SAML) metadata format. A sample of the trust fabric and SAML assertions is contained in Appendix C. From a security perspective, the value of the trust fabric is that it is an electronically signed document that includes the certificates of all trusted systems participating in the MISE. The electronic signing of the trust fabric ensures tampering can be detected and the certificates in the document can be trusted. The certificates are used to establish SSL connections for all other communications.

Once a trusted system has been added to the trust fabric, every connection to the ISI, to publish or consume information, includes information about the trusted system and the user system it's representing. At multiple points during the transaction, the ISI validates the request, the attributes and the data to ensure only the correct information is shared.

Figure 11 shows the steps in a typical transaction. Each time the request or the information touches the "ISI: Policy Enforcement Service" bar, the attributes of the trusted system and the user are compared against the security attributes assigned to the information by the data provider to ensure the transaction is allowed. Throughout the process, every step and decision is entered in the audit log.

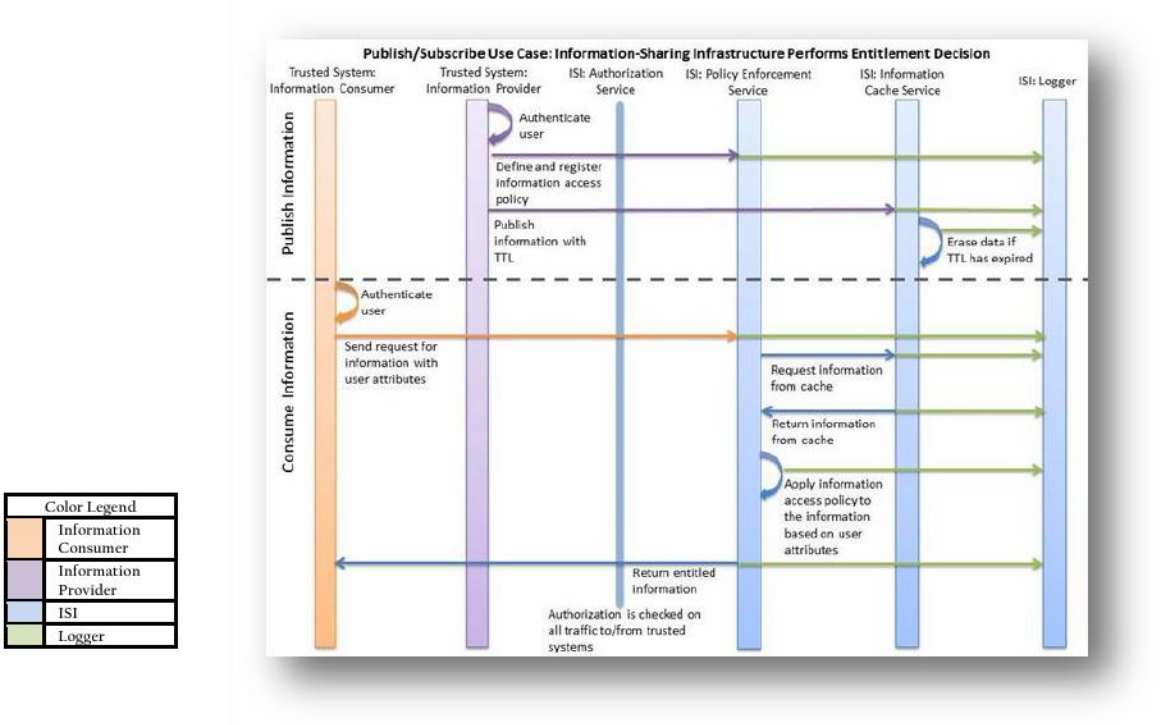

Figure 11: Information Sharing Sequence Diagram

6. Governance

Successful information sharing can only be achieved if information providers can be assured that information is protected in accordance with applicable doctrine. The MISE governance is comprised of three major areas:

1. Service Management
    a. Service Level Agreements
2. Organizational Roles and Responsibilities
    a. MISE Board of Directors
    b. MISE Management
    c. Trusted Systems
3. End User Support
    a. Three Tiers – Local, Trusted System, MISE

Full details of the proposed MISE governance are contained in Appendix F.

The below figure depicts the major areas of the MISE governance.

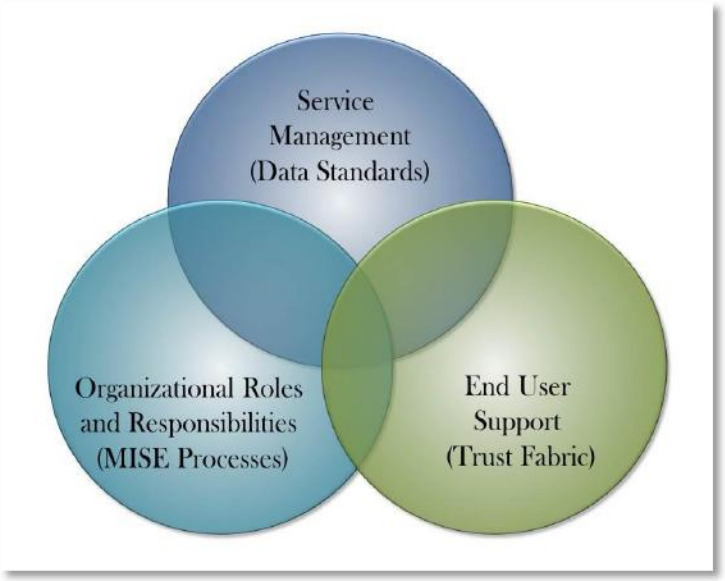

Figure 12: MISE Governance

6.1. SERVICE MANAGEMENT

To correctly govern the information sharing process, the MISE and each trusted system will need to have service level agreements (SLA) which formally define the level of service that should be expected from each. The development and management of MISE SLAs will be the responsibility of the MISE Management Organization. The service level agreements will define:

- The type of services that will be used for the information exchange
- Terms for selecting, accepting, maintaining and transitioning services into the MISE
- Terms for executing and evaluating the quality of service of information exchanges
- Terms for executing the service agreement

Terms for executing the service agreement will include:

- Contacts and Role assignment - Responsible organizations and roles that have access to the trusted systems and the MISE
- Reporting - Logs will be provided based on terms laid out in the SLA
- Review Process - Quarterly reviews will be established that include discussion on SLA fulfillment and future projects that may affect the SLA
- Performance Level Guidelines - Outlines the inbound and outbound information flows to and from the ISI and expected behavior from each trusted system and the MISE
- Uptime Requirements - Outlines the expected availability that the trusted systems must deliver as well as the expected availability that it will provide
- Equipment Support Requirements - Access and security of trusted systems in the MISE

- Third Party Sharing – Can data be shared beyond the receiving trusted system other than as defined in the releaseability attribute?

6.2. ORGANIZATIONAL ROLES AND RESPONSIBILITIES

The roles and responsibilities for the execution of the MISE are divided between three organizational groups.

MISE Board of Directors

The MISE Board of Directors is the executive-level body with representation from primary stakeholders that guides the MISE and is the final authoritative body to make decisions for the environment.

MISE Management Organization

The MISE Management Organization is responsible for the day-to-day operations for the MISE and will facilitate the identification, collaboration, management, movement, and processing of maritime information by leveraging the MISE. The MISE Management Organization is also responsible for the establishment and management of the Configuration Control Board.

Trusted Systems

Trusted systems, both Information Providers / Information Consumers, are responsible for adhering to the information sharing environment rules established by MISE management.

6.3. END USER SUPPORT

The MISE help desk structure focuses on solving problems as close to the user as possible. The MISE Technical Authority serves as the primary point of contact for the MISE Help Desk, but complex technical issues across the environment will be handled according to the following three-tier structure:

Tier 1: Local Help Desk Support

All issues encountered by users should be first reported to the users local help desk. This level of user assistance is provided by the user's local department to resolve simple issues reported by users.

Tier 2: Trusted System Help Desk Support

Any issue the local help desk cannot resolve will be elevated to the Tier 2: Trusted System Help Desk Support level. The trusted system help desk will attempt to resolve the issue, and will contact the MISE help desk if the issue relates to an information provider or the ISI.

Tier 3: MISE Help Desk Support

Any issue that cannot be resolved at the Tier 2 level should be escalated to the Tier 3: MISE Help Desk. Issues that are resolved at this level will be tracked in the MISE Technical Authority issue tracker database.

7. Implementation

Information sharing has become the common thread to success and mission accomplishment for agencies and organizations across all levels of government (Federal, state, local, tribal, and territorial) as well as the private sector. While information sharing can benefit National Security, it must be done consistent with legal and policy constraints. The Architecture Plan defines a low cost, technical solution that is easily implemented. The Architecture Plan also puts mechanisms in place to manage legal and policy concerns. The uniqueness and benefit of the Architecture Plan is that it allows agencies and organizations to utilize their own requirements, authorities, information infrastructures, and resources.

It is not feasible to develop an environment to accommodate specific concerns of every agency; however, the MISE supports sharing and management of the majority of unclassified maritime information. Having this common central hub for unclassified maritime information that the Nation collects and processes will support and improve National Maritime Security.

7.1. IMPACT & SCOPE

The NMDAP and the National Strategy speak to all levels of partnership sharing, public, private, domestic and international. The concepts contained in the Architecture Plan have been reviewed and adopted within the U.S. Government and is being shared with other domestic partners. The NIEM-Maritime standard has been reviewed and is being adopted by domestic *and* international partners. The impact and scope of the Architecture Plan supports the core principles, goals and objectives in the National Maritime Domain Awareness Plan and the National Strategy.

7.2. TRUSTED SYSTEM IT REQUIREMENTS

Becoming a trusted system is the first and most important step to becoming part of the MISE. To be considered for a trusted system and entered into the trust fabric, a system must:

- Have a user management system. The system must manage anticipated and unanticipated (registered and unregistered) users. The system must be able to provide attributes applicable to every user.
- Provide for attribute management through the trusted system for third party sharing.
- Be able to implement RESTful web services that produce and consume NIEM compatible messages.
- Demonstrate the need to have access to the information contained in the MISE.
- The approval of the MISE Governance Board.

7.3. RESOURCE REQUIREMENTS

The "cost" of the Architecture Plan and its resourcing will continue to be a discussion point within the interagency community. Ultimately, the MDA Executive Steering Committee will determine who is responsible for the various parts of the MISE. It will however be the responsibility of the departments, agencies or offices that own and maintain the trusted systems to ensure those systems are properly resourced and maintained.

FRAMEWORK

1. Overview

The National Maritime Domain Awareness Plan describes an information sharing end-state in which an authorized user is able to access the right information at the right time. Access is managed through entitlements based on the needs, rights and authorities of both the user consuming the information and of the organizations providing the information. Although a key MDA tenet is to provide the broadest information sharing possible, it can only be achieved if information providers are assured the information is protected according to their requirements. The establishment of the Maritime Information Sharing Environment (MISE) and the data standards as described in this plan is the first step toward that assurance.

The following figure depicts the Operational Overview (OV-1) for an information sharing environment as a national capability to share and search maritime information across organizational boundaries. Within the information sharing environment depicted in Figure 13, an infrastructure exists to effectively link participants and information. This results in a distributed, protected, and trusted environment.

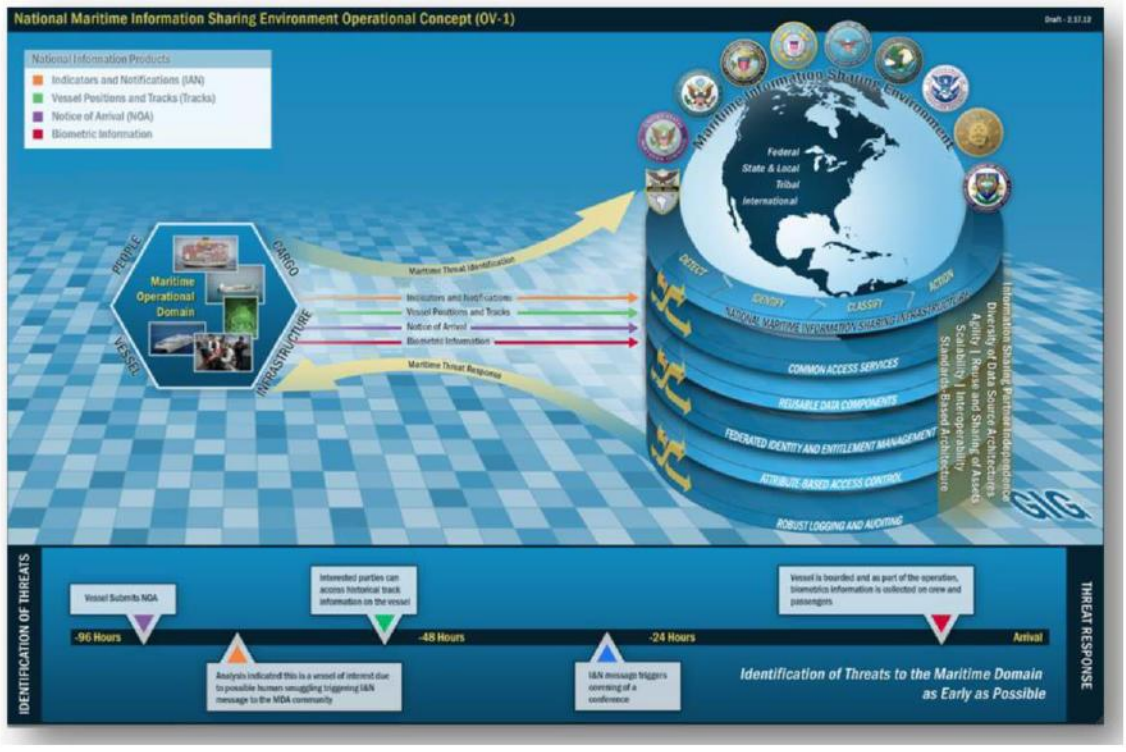

Figure 13: MISE Operational View

The Operational Overview identifies four subject areas that make up the Maritime Operational Domain: Cargo, Infrastructure, Vessels, and People. Each of the participants in the Global

Maritime Community of Interest (GMCOI) contributes to the National MDA in terms of these four subject areas. Specifically, the GMCOI interacts regarding the monitoring and collection of this information. The MISE focuses initial efforts on securely sharing the following *National* information products within the GMCOI:

1. Notice of Arrival (NOA) - A 96-hour advance notice that vessels over 300 gross tons and inbound to US ports are required to submit which lists vessel, crew, passenger, and cargo information.
2. Indicators and Notifications (IAN) – (including warnings and alerts):
    a. Indicator – Any piece of discreet information, processed or otherwise that is used to inform or contribute to an analytical process
    b. Notification – Directed communication
        i. Warning – Notification of a possible activity or event
        ii. Alert – Notification of the actual or eminent execution of an event.
3. Positions (POS) - A geospatial position, course, heading, speed, and status of a vessel at a given time. A series of position reports can be combined to produce track information.

2. Objective

The overarching goal of the Architecture Plan is to establish an information sharing process and framework to advance information sharing capabilities. The objective is for agencies and organizations within the GMCOI to be able to utilize their own requirements, authorities, information infrastructures, and resources to share information within the community that would assist operational missions. The MISE simply provides an environment for the successful transfer of information between information consumers and providers.

This section provides technical managers with greater detail for the technical implementation of the architecture. Documents and processes that are required to conduct specific actions within the architecture are referenced, but not explained in detail.

3. Definitions

The definitions below are used in the following sections to describe the participants and processes within the MISE.

Attributes	Characteristics of a persona that defines a user or system in a particular role (sent by a trusted system to the ISI).
Decision Authority	The authoritative role to enforce information access policy based on user attributes to determine what information will be provided to a user.
Entitlement(s)	The result of determining access by evaluating IAP against user attributes.

Identity Provider	A service provider that creates, maintains, and manages identity information and provides authentication services to other systems.
Information Access Policy (IAP)	A rule set that defines the attributes users must possess to allow access to information.
Information Consumer	A type of trusted system that authenticates users, using an internal or external identity provider, and passes information access requests and user attributes on behalf of the user to the ISI.
Information Provider	A type of trusted system that provides both maritime information and corresponding security attributes to the ISI.
Information-Sharing Infrastructure (ISI)	Handles requests from the trusted systems on behalf of the users. Operating on the user attributes, it will make entitlement decisions, processing messages from the information providers based on their security attributes and providing responses to the trusted system.
Request Broker	A role of the ISI to forward requests for information from the information consumer to the information provider.
Service Provider	An entity that provides services to other entities.
Trusted System	A system that is recognized and authenticated by the ISI and operates in a role of information provider and/or information consumer. It also passes information or information access requests and user attributes to the ISI. The trusted system relies on its own internal processes to authenticate users.
User	Any person, organization, system, etc. that authenticates via the trusted system in order to access the ISI.

Table 1: MISE Term Definitions

4. Information Sharing Process

Data movement within the MISE is between trusted systems and the ISI. Central to the information sharing process are two points. First, the decision authority either rests within the ISI or the information provider. That is, either the ISI or the information provider compares the IAP with the user attributes and decides what information should be provided back to the user. Second, information will either be cached within the ISI or stored by the information provider. Caching the information, allows the ISI to quickly respond to requests and to protect the information provider systems from query overload. These two points are addressed within the following two use cases of the ISI. Note that in all use cases, the same IAP is applied.

In use case 1, the ISI caches the information *and* acts as the decision authority.

In use case 2, the ISI acts as a request broker, passing both the request and the user attributes from the information consumer system to the information provider system. The information provider system acts as the decision authority.

The following table summarizes roles with respect to each of the use cases. The distinguishing characteristic between the use cases pertain to which system has decision authority and whether or not caching is employed.

Function	Use Case 1	Use Case 2
Caching	ISI	N/A
Decision Authority	ISI	Information Provider
Logging	ISI	ISI
Request Broker	N/A	ISI
Trusted System Authentication	ISI	ISI
User Authentication	Trusted System	Trusted System

Table 2: Roles Across Each Use Case

4.1. USE CASE 1: ISI ACTS AS DECISION AUTHORITY AND PERFORMS CACHING

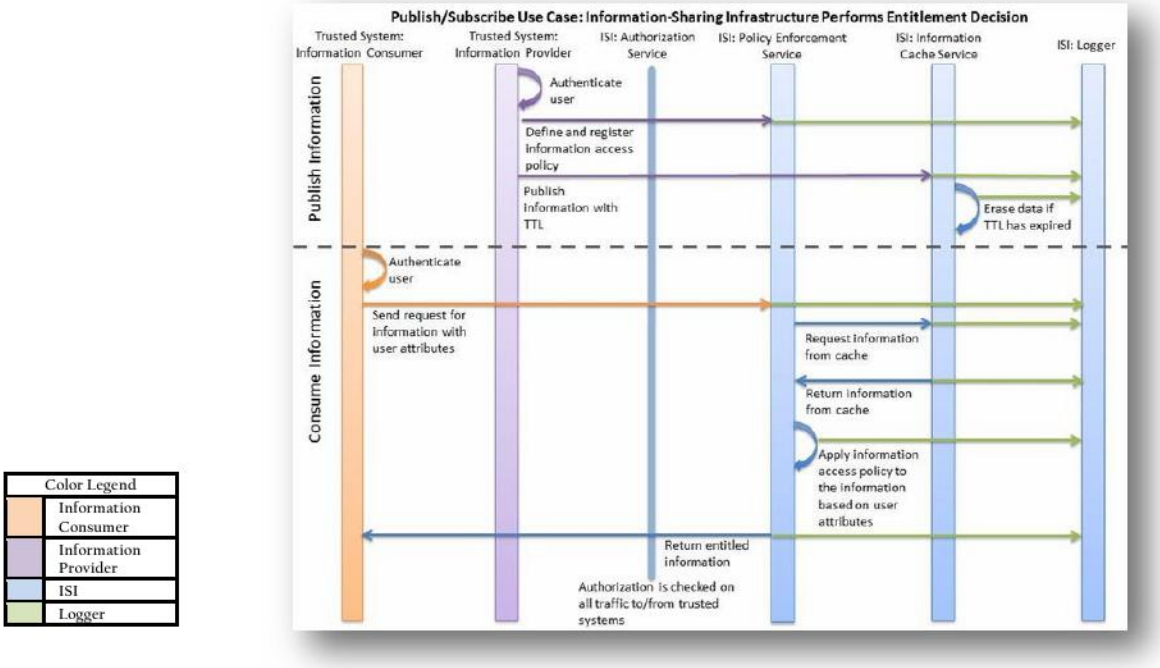

Figure 14: Use Case 1 Process

4.1.1. PUBLISH INFORMATION

A trusted system, in the role of an Information Provider, authenticates a user. The Information Provider can publish information with a Time to Live (TTL) to the ISI's information cache. This allows for messages to be published by the information provider on a constant basis and cached by the ISI to handle repeated queries from consumers for the same information. Any information in the ISI cache is purged when the TTL is reached.

4.1.2. CONSUME INFORMATION

A trusted system in the role of an Information Consumer authenticates a user using either an internal or external identity provider. From that trusted system, the user makes a request for information from the ISI. The trusted system supplies a set of attributes that define that user's persona along with the request to the ISI. The ISI queries its own local cache for matching information and compares the IAP attributes in each record from the data provider to the user attributes provided with the request. By evaluating the IAP against the user attributes, the ISI makes an entitlement decision. According to the IAP, the message response is tailored to match the entitlements and the resulting information is returned to the user.

4.2. USE CASE 2: ISI ACTS AS REQUEST BROKER, INFO PROVIDER ACTS AS DECISION AUTHORITY; NO CACHING

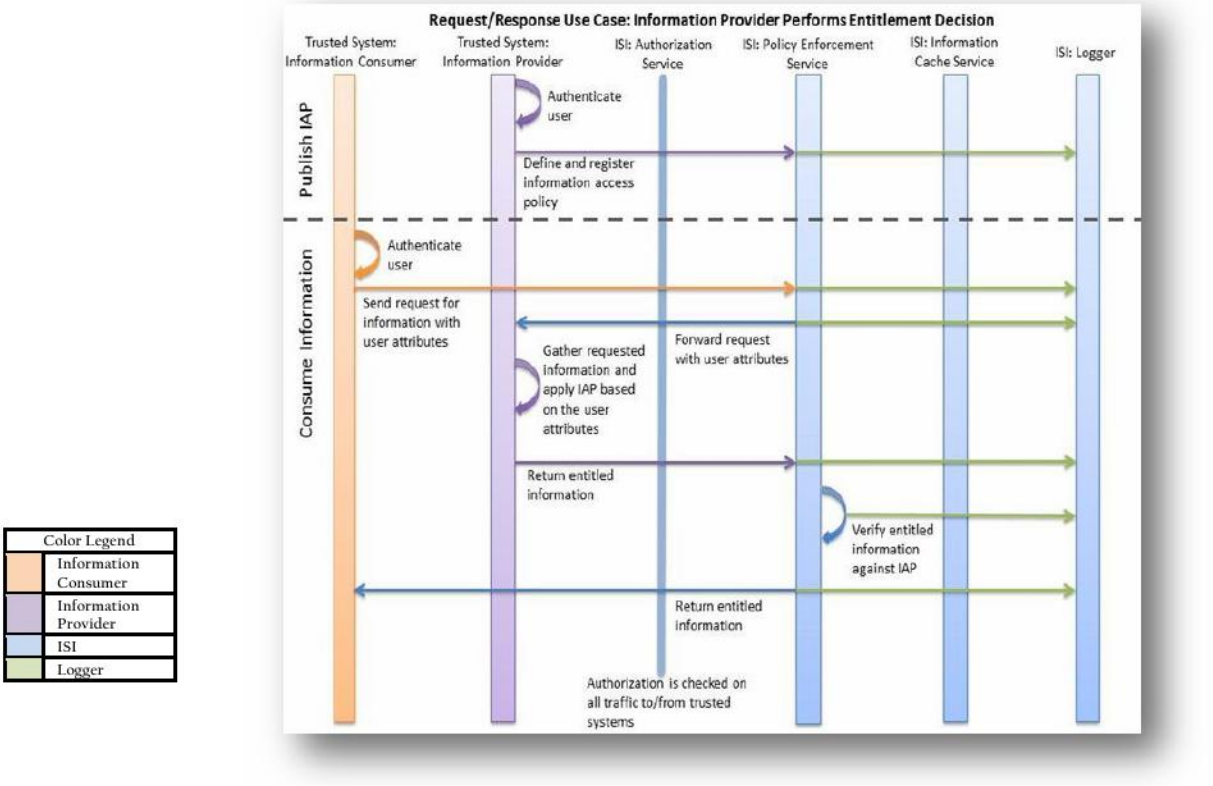

Figure 15: Use Case 2 Process

4.2.1. PUBLISH INFORMATION

A trusted system in the role of an Information Provider authenticates a user. Since the trusted system is authorizing the user request and providing the requested data, no data is published to the ISI.

4.2.2. CONSUME INFORMATION

A trusted system in the role of an Information Consumer authenticates a user. From that trusted system, the user makes a request for information from the ISI. The trusted system supplies a set of attributes that define that user's persona. The ISI forwards the request with user attributes to the information provider. The information provider gathers the requested information and applies their IAP based on the user attributes. The information provider returns the entitled information back to the ISI. The ISI verifies the entitled information against the IAP and returns the resulting information to the user.

4.3. COMPARING THE USE CASES

Information providers will implement the use case that best meets their requirements. Table 3 provides a comparison of the two proposed use cases from the information provider's perspective.

Use Case	Benefits	Implications
Use Case 1 ISI Acts as Information Broker (ISI has Decision Authority and ISI Performs Caching)	Eliminates cost of servicing requests Reduces implementation costs associated with enforcing information access policy	Information Provider must tag messages with security attributes to convey information controls Information Provider must keep cached information in the ISI current
Use Case 2 ISI acts as Request Broker only (Information Provider has Decision Authority; No Caching)	Information Provider maintains maximum control over data by assuming IAP enforcement responsibility	Increased implementation costs associated with enforcing information access policy Information Provider must tag messages with security attributes to convey information controls Information Provider must ensure capacity to service requests from National Maritime Federation

Table 3: Use Case Comparison

4.4. INFORMATION SHARING INFRASTRUCTURE (ISI) FEATURES

The Information Sharing Infrastructure (ISI) is the central hub of the MISE and the component that provides the service interfaces to the trusted systems. At a high level, the ISI features are captured in three functional areas: Entitlement Management, Information Management, and Logging/Auditing.

The features of the functional areas are:

Entitlement Management

ENT-001. Users are granted access to maritime information based on information access policies and attributes assigned by the information providers.

ENT-002. The ISI recognizes a standardized set of security attributes for entitlement management as defined in the National Architecture Plan. Information providers use the security attributes to tag information they publish to the ISI with metadata to convey data protection requirements. The security attributes are included as metadata on the information provided to consumer systems for integrity throughout the data management lifecycle.

ENT-003. Information providers can apply security attributes at varying levels of granularity depending on policy: message-level (e.g., an entire XML message); record-level (e.g., records within an XML message).

ENT-004. Trusted systems authenticate users using either an internal or external identity management process. Trusted systems supply user attributes for all queries. The user attributes and entity attributes (corresponding to the trusted system) are compared to the information provider supplied security attributes on the data to make entitlement decisions.

ENT-005. Users are prohibited from accessing information for which they or their trusted system do not have entitlements.

ENT-006. Both information provider and information consumer trusted systems are authenticated by the ISI. All connections between the ISI and the trusted systems will be secured via SSL.

ENT-007. Trusted systems will be responsible for complying with all rules for maintaining trusted system status.

Information Management

INF-001. Ingest, index, and cache messages from appropriate information providers.

INF-002. Information providers can set the expiration date of their messages, which shall not exceed 30 days from the time of publication; the ISI will automatically remove expired messages from the cache when they expire.

INF-003. Search and retrieve information from the cache to service requests from trusted systems.

INF-004. Securely dispose of expired information.

Logger & Auditing

LOG-001. Log all external interactions with the ISI.

LOG-002. Log all internal system activities.

LOG-003. Log all attributes presented by a trusted system for each request, the request itself, the messages retrieved to match the request, and the resulting messages returned to the trusted system upon an entitlement decision.

LOG-004. Log all attempted and established connections and sessions.

LOG-005. Capture all states and activities necessary for auditing.

LOG-006. Log all message disposal activity.

LOG-007. Maintain all logs in accordance with all auditing requirements. Ensure that logs are fully backed up for auditing purposes.

LOG-008. Log all auditing activity.

5. Maritime Information Sharing Environment

The processes required for the GMCOI to both share and protect information are explained in this section. There are key components of the MISE that are critical to understanding before explaining how a member of the GMCOI becomes a participant or trusted system. The figure below is a physical view of the systems and roles in the MISE.

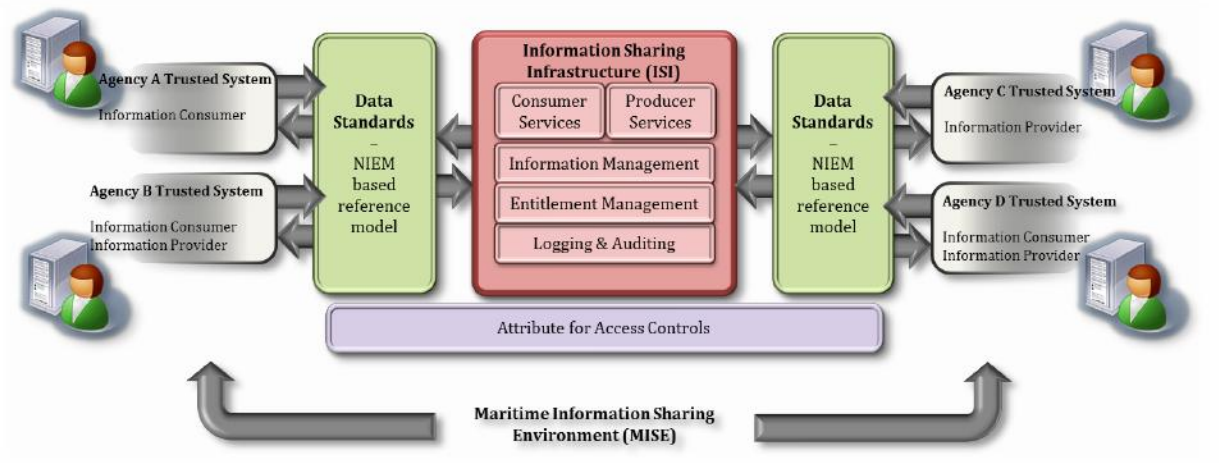

Figure 16: MISE Physical View

The *MISE Physical View* illustrates several key characteristics of the MISE, and introduces some terminology:

- The MISE consists of many separate systems that are owned and operated by separate agencies. These systems and agencies work together in a *federation*. In most cases these are existing systems, already fulfilling key roles within their respective agencies. To participate in the MISE, they will be expected to adhere to MISE specifications, allowing them to share information within or retrieve information from the information sharing environment.
- The *Information Sharing Infrastructure (ISI)* is a system operated by MISE Management, and acts as the central hub of the environment. The ISI provides essential services to the systems within the MISE. The role of *MISE Management* is discussed in the appendices.
- Each system within the MISE is referred to as a *Trusted System*.
- To participate in the MISE, trusted systems connect to and interact with the ISI, not directly with other trusted systems. Network connections between trusted systems and the ISI are over the public Internet.

- *Information Provider* and *Information Consumer* are the two key roles which a trusted system can take within the MISE. Any trusted system may serve as either or both of these roles, as illustrated in the figure *MISE Physical View*.
    - An *Information Provider System* makes information available to other trusted systems, and maintains control over which permissions users must possess to access the information. The information provider system controls access by providing security attributes to the ISI for each record published.
    - An *Information Consumer System* retrieves information from the MISE. The consumer system is given access to information within the MISE based upon the applicable *Information Access Policy*.

Users in the MISE are associated with a specific trusted system which conducts interactions such as authenticating the user and retrieving information from the ISI on behalf of the user. This trusted system is the conduit for the user to access information from the ISI or other systems in the MISE.

6. National Maritime Architecture Framework

Maritime information must be integrated and shared across numerous stakeholder agencies to gain an effective understanding of the maritime domain. To allow sharing and integration across organizational boundaries, information providers must be assured that their information is protected with access granted only to those parties entitled to the information in accordance with applicable laws and regulations. The National Maritime Architecture Framework is an interoperable solution and approach to enable stakeholders across the GMCOI to share and protect their information.

The National Maritime Architecture Framework is comprised of four architectural views:

1. Data Architecture View
2. Services Architecture View
3. Security Architecture View
4. Technical Operations Architecture View

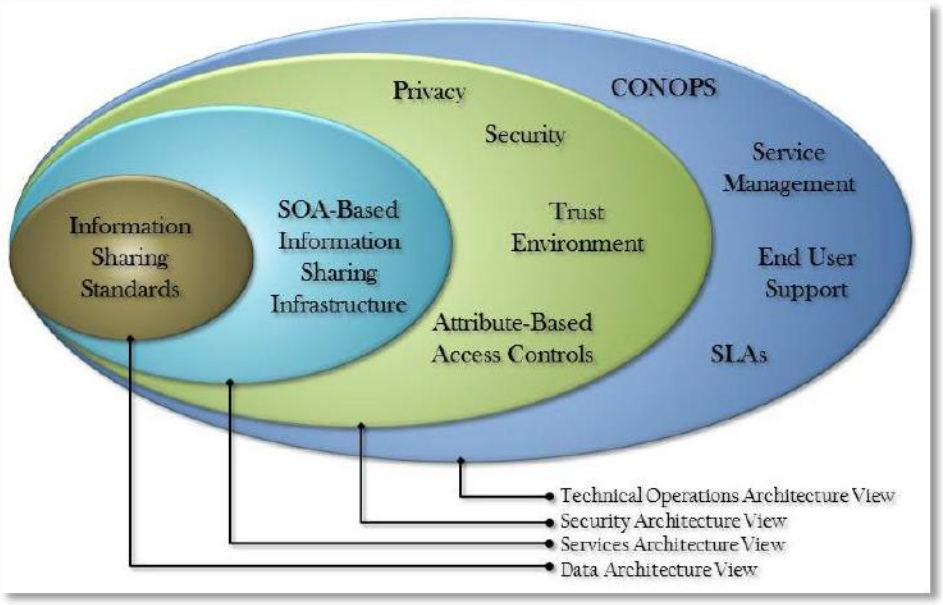

Figure 17: The Four Architectural Views of the Framework

Data Architecture View

Shared vocabulary and information sharing standards based on the National Information Exchange Model – Maritime (NIEM-M) reference models form the foundation of the framework. In accordance with the National MDA Architecture Hub Strategic Plan, the Maritime Framework leverages NIEM for a common vocabulary and consistent processes to build maritime information exchange standards.

Services Architecture View

The information sharing infrastructure (ISI) is the core functionality for the Maritime Information Sharing Environment. The ISI provides a common integration platform to facilitate management and dissemination of maritime information across the maritime community. The categories of services within this platform that are provided by the ISI include: Information Brokering, Request Brokering, Logging, Auditing, Information Management, and Entitlements Management.

Security Architecture View

The Trust Environment employs standards-based Federated Identity Management for secure, seamless access to maritime information products, which are then shared among trusted systems and the ISI. The Trust Environment consists of attribute-based access control, common user attributes/entitlements metadata, SAML[11] profile for entitlement assertions, and fine-grained, information provider managed access control policy. This environment ensures only trusted systems interact with the ISI and that user entitlements are determined based on user attributes.

[11] Security Assertion Markup Language (SAML) is an XML-based open standard for exchanging authentication and authorization data between systems.

Technical Operations Architecture View

The Technical Operations Architecture View defines how the framework as a whole is managed and operated. It includes governance and agreements used to manage the MISE.

6.1. DATA ARCHITECTURE VIEW

6.1.1. INFORMATION MODEL

While initial focus is on three national information products[12], a guiding design principle allows other types of information to be handled without requiring modification to participating trusted systems.

The *information model* consists of a set of characteristics that are true of all shared information. These characteristics are:

- A *record* is the unit of shared information. Records are represented as XML documents.
- Each record has a *record type*, which is defined by a NIEM IEPD (discussed below). The record type can be equated to a record subject; Vessel, People, Notification, etc.
- A *Data Set* is a set of records shared into the environment by an information provider. All records within a data set have the same record type. It is possible for an information provider to share more than one data set, if they have multiple types of records to share. Conversely, multiple information providers may share data sets of the same record type.
- An *Information Access Policy* defines the access to records within the data set. Information providers convey information access policy by applying security attributes at the record logical block, or element level in the records they publish.

The ISI and information consumer systems define and build services based on knowledge of this information model and record type definitions, while relying on the consistent representation of information provided by the NIEM-M vocabulary. Specialized functionality can and will be built by information consumer systems based on detailed understanding of specific record types. The information model concept is unique because it allows a base level of functionality to operate upon record types defined and introduced after the information consumer system is implemented. The concept is similar to polymorphism in object oriented programming, i.e. code written to operate on objects of a base type can still deliver some useful functionality on objects of a more specialized derived type.

6.1.2. NATIONAL INFORMATION EXCHANGE MODEL

The National Information Exchange Model (NIEM) is an XML-based information exchange framework for sharing data between United States government agencies and their information partners.

The NIEM facilitates the creation of automated enterprise-wide information exchanges, which can be uniformly developed, centrally maintained, quickly identified and discovered, and efficiently reused. The result aims toward more efficient and expansive information sharing

[12] Notice of Arrival (NOA), Indicators and Notifications (IAN), and Positions and Tracks.

between agencies and jurisdictions; more cost-effective development and deployment of information systems; improved operations; better quality decision making as a result of more timely, accurate, and complete standardized information; and, as a consequence, enhanced public safety and homeland security.

The NIEM framework has several components:

- **NIEM Core**: A common XML-based data model that provides data components for describing universal objects such as people, locations, activities, and organizations.
- **Domains**: More specialized XML data models for individual use cases. There is a specialized domain for Maritime, along with a number of others (e.g. Justice and Immigration).
- **Information Exchange Package**: A methodology for using and extending the building blocks that come from the common and domain-specific models, and turning them into a complete information exchange.
- **Tools**: Help develop, validate, document, and share the information exchange packages.
- **Governance Organization**: Provides training and support and oversees NIEM's evolution over time.

The NIEM-M XML vocabulary provides a combination of objects from NIEM core, the Maritime domain, and additions via the EIEM and IEPD, described in the next sections.

6.2. ENTERPRISE INFORMATION EXCHANGE MODEL

An Enterprise Information Exchange Model (EIEM) is a NIEM-conformant set of XML schemas and other artifacts. The EIEM defines core entities, also termed Business Information Exchange Components (BIECs) that serve as building blocks that are reused in many maritime exchanges.

The high-level BIECs defined by the Maritime EIEM are:

- CDC Cargo (Certain dangerous cargo as declared in an Advance Notice of Arrival)
- Crew Nationality Count
- Interest
- Movement
- Non Crew Nationality Count
- Port Visit
- Position
- Vessel Characteristics
- Vessel History
- Vessel Identification
- Voyage Information
- Record Metadata

Using an EIEM increases the commonalities among models, which has a number of benefits:

1. Simplifies development and allows users to develop reusable processes that work across multiple record types.
2. Reduces complexity by reducing the number of overall elements that need to be supported.
3. Decreases the time and cost of maintaining the model because it is simpler and has less overlap.

6.3. INFORMATION EXCHANGE PACKAGE DOCUMENTATION

An Information Exchange Package Documentation (IEPD) defines a particular record type, the basic unit of shared information. The first IEPDs for the initial focus areas are:

- Advance Notice of Arrival (NOA)
- Indicators and Notifications (IAN)
- Vessel Positions (POS)

As depicted in the figure below, an IEPD uses a subset of the core BIECs defined in the EIEM and assembles them into a particular record type with its own unique root element. As new requirements are defined, new IEPDs can be created that build on the same EIEM core entities.

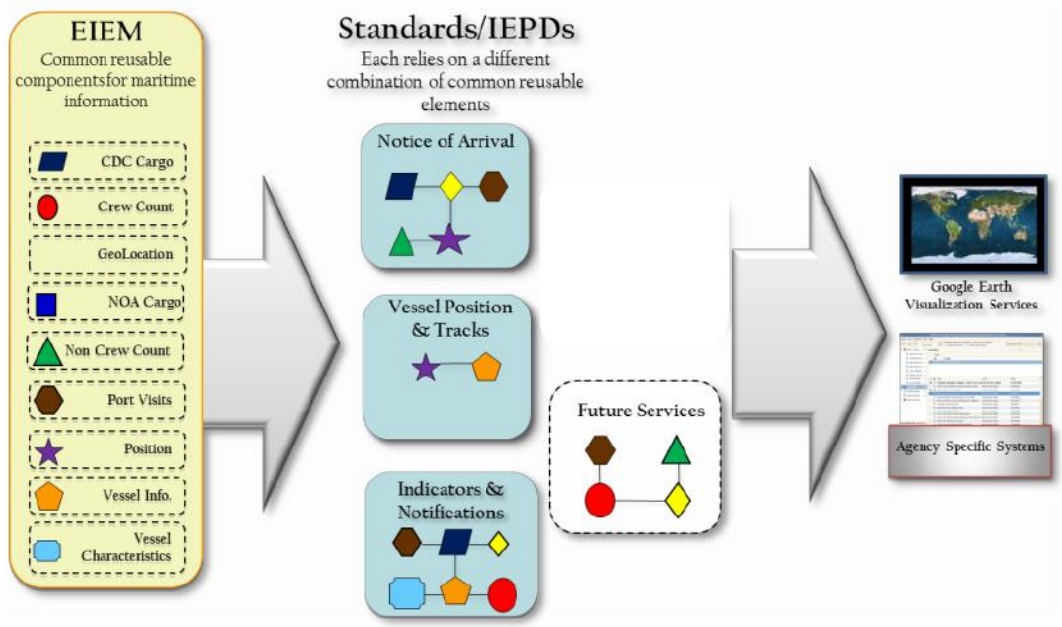

Figure 18: IEPDs combine core entities from the EIEM

The MISE is flexible enough to allow an IEPD to extend or restrict the EIEM components, as long as no required properties are removed. An IEPD can also define new exchange-specific information, for example a notice date for a notice of arrival.

Physically, an IEPD is a collection of artifacts that describe the exchanged data, including:

- XML Schemas that describe the structure of the XML documents:
  - *Exchange schema* that declares the root element of the exchange
  - *Extension schema* that contains exchange-specific components
  - Copy of the EIEM schema that defines the BIECs
- Documentation artifacts such as a schema description and change log
- Metadata artifacts such as rendering instructions, and information to aid query building
- Sample XML instances

6.4. SERVICES ARCHITECTURE VIEW

The Architecture Plan defines the MISE service interfaces that allow trusted systems to retrieve and publish information. Trusted systems interact with the ISI using these service interfaces.

6.4.1. SERVICE TYPES

RETRIEVING DATA

Two services are defined for retrieving information:

- **Search Service**: Find records matching specified criteria. Returns a list of summary records and record IDs in an atom feed.
- **Retrieval Service**: Retrieve a record with a specified ID. Returns the NIEM-M representation of the record.

These services provide access to records, which are the unit of information the MISE operates upon. Information consumer systems may choose to present a relatively low level user interface allowing search, retrieval, and viewing of individual records. Information consumer systems may also use information obtained from multiple service invocations to compose a higher-level operational picture for presentation to users.

PUBLISHING DATA

In addition, a single service is defined for *publishing* information into the ISI.

- **Publication Service**: Place information into a cache, and keep the cache up to date including delete operations to remove information from the cache.

Each of these services can be thought of as having a *provider* side and a *consumer* side. Providers *implement* the service and make information available. Consumers *use* the service to obtain information.

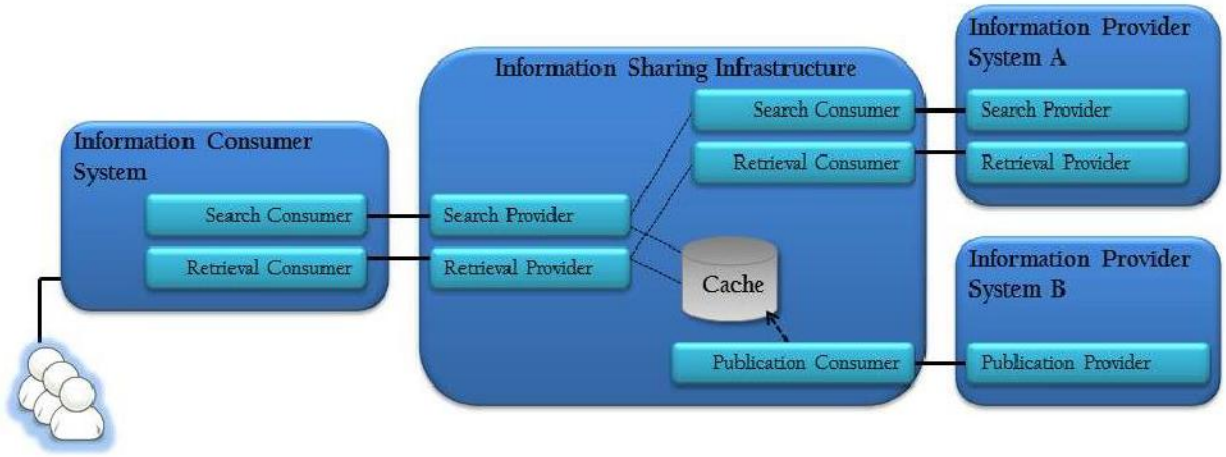

Figure 19: Service Providers and Consumers

The figure above illustrates several important aspects of MISE services:

- From a services perspective, information providers have two integration choices:
  - ***Information* Broker**: Information Provider B (figure above, bottom right) has chosen to publish information to the ISI cache, so it implements only the provider side of the Publication service interface. The ISI is able to provide Search and Retrieval services to information consumers using the cached information. In this case, the ISI acts as an *information broker*.
  - ***Request* Broker**: Information Provider A (figure above, top right) has chosen not to cache information in the ISI. In this case, each time an information consumer system searches the ISI using the Search service, the ISI must in turn search the information provider system. The same is true for Retrieval. So in this case, the information provider system must implement the provider side of the Search and Retrieval services, and the ISI implements *federated search* and *brokered retrieval*. In this case, the ISI acts as a *request broker*.
- The same service interfaces are used in different contexts. For example, information provider systems that do not use the ISI cache provide the same Search interface to the ISI as the ISI provides to information consumer systems. Depending on the role of a trusted system (information provider and/or consumer), and its integration choices, provider and/or consumer sides of several MISE service interfaces may need to be implemented.

The ISI serves as the common integration point, allowing each trusted system to make integration choices, while still enabling interoperability.

6.4.2. SERVICE CHARACTERISTICS

This section describes characteristics that are common across all service interactions between trusted systems and the ISI. Subsequent sections provide more details of each service.

NETWORK CONNECTIVITY

Trusted systems connect to the ISI over the public Internet. All service interactions use the HTTPS protocol.

RESTFUL

MISE services are RESTful web services[13]. RESTful web services use the HTTP protocol as an application protocol, rather than simply as a transport. RESTful web services are implemented very much like standard web applications, as opposed to using any type of complex middleware layer. The RESTful style was chosen for MISE because of its simplicity and ease of integration.

As RESTful services, nearly all MISE service interactions follow a simple pattern, which is the same pattern used by a web browser accessing a page on the Internet:

- The service consumer requests an HTTP GET of a URL specifying a resource at the service provider. The URL is known to the service consumer because it is one of the following:
    - A top-level URL of the service provider obtained from configuration information.
    - Relative to the top-level URL, as specified in the MISE Search/Retrieve Interface Specification.
    - Obtained from the body of an earlier response from the service provider.
- The service provider returns a resource in response to the GET. The format of the resource is specified in the MISE Search/Retrieve Interface Specification. In most cases it is an XML document adhering to a specified schema or format.
- Any error conditions are reported using standard HTTP status codes.
- Other aspects of the exchange (e.g. session handling, conditional get, content negotiation) are handled following standard HTTP conventions.

AUTHENTICATION AND ENTITLEMENT

All MISE services operate in the context of the security architecture discussed in section 6.5. For every service interaction, provider and consumer systems are authenticated. If a service invocation occurs on behalf of a user, authenticated attributes of the user are available. From the standpoint of service definitions, the security architecture operates conceptually when the connection is established, not as a part of each HTTP request/response message.

6.4.3. SEARCH SERVICE

The Search service allows a consumer to locate records, which match specified criteria. The Search service provides a flexible query language. In most cases, the search service does not return full records, but instead returns the IDs of matching records. If a consumer wishes to obtain a full record, an ID obtained from the search service is used with the Retrieval service to obtain it.

The core of the Search service interface consists of:

[13] The term "REST" stands for REpresentational State Transfer and is an architecture style for designing networked applications.

- The Search consumer requests an HTTP GET of a URL such as https://*providerbase*/search?q=*query*, where *providerbase* is a base URL obtained from configuration information, and *query* provides search criteria expressed using the Search service query language, which is specified in detail in the MISE Search/Retrieve Interface Specification. The query language allows searching for records containing elements from the NIEM-M vocabulary with specified values. It also allows constraining results to specific record types, providers, and data sets.
- The Search provider responds with an Atom document (which is an XML document following a specific schema), listing IDs (which take the form of URLs) of matching records, and other information described in the MISE Search/Retrieve Interface Specification. To be included in the list, the portion of the record that the user is entitled to see (according to the IAP) must match the search criteria. If a large number of records match the search criteria, only a subset is included to constrain the size of the response. In this case, the response also includes the total match count, and a URL to the next page of results.

Information consumer systems often have a need to monitor the MISE for new information matching specified criteria. This capability is provided using the *web syndication* model. The information consumer system periodically queries the Search service, using the full search query language, but including additional parameters allowing the returned Atom document to report recent changes in a manner efficient for the provider and consumer.

The ISI also supports some specially defined *value-added* searches, allowing related information from a number of records to be combined, and returning full information in response to the search request. In this case, it is not necessary to use the Retrieval service.

Another part of the Search service interface specifies relative URLs to metadata describing data sets, information providers, and record types, which are available at the Search provider. This information is used by Search consumers to form proper queries and to allow building of a search user interface.

6.4.4. RETRIEVAL SERVICE

The Retrieval service obtains the full NIEM-M representation of a specified record. The Retrieval service interface consists of:

- Retrieval consumer requests an HTTP GET of a URL representing a specific record. This URL will have been obtained using some previous service invocation, for example from Search results.
- Retrieval provider responds with the NIEM-M representation of the specified record. The Retrieval provider ensures that the IAP is applied so consumers are provided information at their level of access.

An information consumer system may use the Retrieval service for the purpose of building a higher-level operational picture, in which case it may not be necessary to display the individual record to a user. In other cases, an information consumer system may display the record in a manner appropriate and intuitive for individual viewing. Metadata associated with the record type, containing instructions for converting the NIEM-M representation to HTML, can be read by information consumer systems.

6.4.5. PUBLICATION SERVICE

The Publication service allows an information provider to maintain an up-to-date copy of shared information at another location within the MISE. The Publication service interface is used by information providers to update the ISI cache, enabling the ISI to provide Search and Retrieval services on their behalf. In this case, the ISI is the Publication *consumer*, and the information provider is the Publication *provider*.

A Publication provider makes three categories of resources available, which can be accessed by a Publication consumer using HTTP GET:

1. **Metadata** – lists available data sets and record types.
2. **List of all shared records** – resource(s) in Atom format listing IDs of all shared records in the data set.
3. **Ordered list of changes** – resource(s) in Atom format listing all changes to the data set (new records, deleted records, changed records).

These categories of resources provide all information necessary to allow the Publication consumer to initially obtain a full copy of the shared data set, and also to efficiently check for changes. The Publication consumer reads these resources using HTTP GET.

6.4.6. HOW THE SERVICES WORK TOGETHER

The following serves as an example of how the MISE services can be used by information consumer systems to deliver value and insight to users and information consumer system graphical displays showing incoming vessels within a region surrounding a port. Vessels that are mentioned in an IAN record are overlaid with an icon to make this visually apparent. Users of the system can click on a vessel and drill down to see detailed information about the vessel. Users can also choose to display tracks showing the path taken by the vessel to reach its current position.

To implement this functionality, the information consumer system would interact with the MISE in the following manner to obtain information needed:

- A key aspect of implementing this functionality is creating and maintaining the collection of information needed to generate the graphic. This information would be initially obtained using the Search service with a query requesting all records of type NOA, IAN, or POS, with location inside the area of interest (region surrounding specified port). The system would then use the Retrieval service to obtain full records for all matches listed in the search response.
- In order to keep the graphic up-to-date, the system must keep its internal representation of the above information up-to-date. To do so, it periodically uses the Search service with the same query as above, but with additional parameters allowing efficient detection and reporting of changes in the returned Atom document.
- Since the information consumer system wishes to display this graphic to many users, without the overhead of separately searching on behalf of each user, service requests are associated with a set of user attributes representing a group of users. The information consumer system is responsible for only displaying the graphic to individual users who possess these same attributes. If the system wishes to also display a version of the

graphic generated using more privileged information to more privileged users, a separate Search can be requested with a more privileged set of user attributes to maintain this separate internal representation.

- If a user chooses to display the track showing the path taken by a vessel, the system will have the updated position information available in its internal representation.
- If a user clicks on a vessel to drill down and see all available information, the system can show information from the internal representation (e.g. NOA) immediately. The system should also at that point make a Search request for all records of *any* type mentioning the vessel. In this manner, as the MISE expands to include additional types of information (record types), the information consumer system will be able to display this information to authorized users *without the need to modify each information consumer system when new record types are introduced*.
- If an information consumer system was developed with a built-in knowledge of specific IEPDs for known record types, it can display information from those records in any manner it sees fit. To display a record for a user *without* having a built-in knowledge of the record type, the information consumer system can read the metadata associated with the record type from the ISI to obtain instructions allowing generation of HTML presenting the record contents in a manner meaningful to a human viewer.

6.5. SECURITY ARCHITECTURE VIEW

The purpose of the MISE Security Architecture is to ensure that shared information is protected in accordance with applicable doctrine. Information providers must be assured that shared records will only be revealed to users who meet certain criteria. In addition, some records contain subsets of information that should only be revealed to a subset of MISE users. For example, certain portions of records containing privacy-protected information must only be revealed to users who have been authorized to view such information.

This section describes the access control system that meets these requirements.

6.5.1. FUNCTIONAL PERSPECTIVE

We begin by describing the MISE *attribute-based access control system* from a logical or functional perspective. The *Trust Implementation Perspective* section below focuses on how it is secured.

The figure below illustrates key functional aspects of MISE access control.

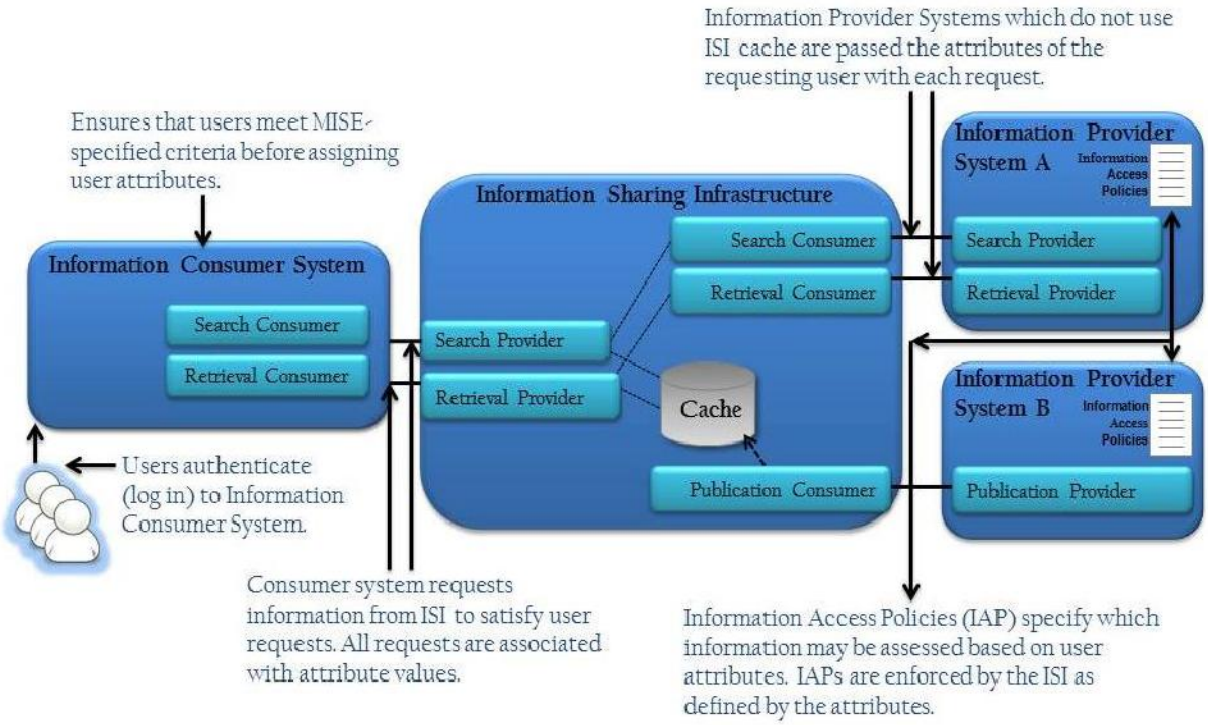

Figure 20: Security Architecture - Functional Perspective

USER ATTRIBUTES

As discussed in section ***6.4.1 Service Types***, information provider systems share data sets with the MISE, and information consumer systems provide users access to the shared information. Users authenticate (log in) to their information consumer system, which then obtains records from the ISI on behalf of the user using the services discussed in section ***6.4.1 Service Types***. Users authenticated by many different consumer systems access information shared by many different provider systems. Providers must have a way to control which users can access which portions of shared information, but it is not practical for providers to possess knowledge of individual users, or even of individual information consumer systems. Rather, a standardized way of describing users is necessary. This is the purpose of standardized *User Attributes*.

The MISE defines a common set of user attributes. Information providers convey *Information Access Policies* by applying security attributes to the records they publish specifying access is allowed for users who possess certain of these attributes, rather than to specific users or to specific information consumer systems. When obtaining information on behalf of a user, information consumer systems associate the set of attributes possessed by the user with the request. Information consumer systems are responsible for user vetting – i.e., ensuring that the attributes for a user are accurate and meet the precise meaning of the attribute as defined by the MISE. Information consumer systems are also responsible for authenticating users, usually by requiring entry of a password or by requiring possession of a token or key.

Specifically, the MISE defines the following user attributes for use by information access policies, referred to as *entitlement user attributes*:

- **Citizenship Code** – The country that has assigned rights, duties, and privileges to the user because of the birth or naturalization of the user in that country.
- **Law Enforcement Indicator** – User required and qualifies for law enforcement information.
- **Privacy Protected Indicator** – User requires and qualifies for privacy-protected information.
- **Community Of Interest Indicator** – Minimum access level assigned; user requires and qualifies for information shared by MISE community.

Other user attributes are defined for auditing purposes, referred to as *auditing user attributes*. They are not used by information access policies, but are recorded in audit logs as a record of the user who accessed information. The auditing user attributes listed below are mandatory. Additional optional attributes are defined in the MISE Attributes Specification.

- **Full Name** – The complete name of the user.
- **Electronic Identity Id** – The unique identifier that is associated with the user within the user's information consumer system.

INFORMATION ACCESS POLICIES

Each shared data set has an associated *Information Access Policy (IAP)*. The IAP is specific to each data set and is maintained by the information provider. An IAP is a rule set that defines attributes users must possess to access each logical block within records in the data set. IAPs allow denying access to entire records, or to portions of records. Rules may be conditional based on values within the record. Note that IAPs apply to all information consumer services (Search and Retrieval). During Retrieval, access to entire records may be denied, or unauthorized portions of records may be excluded from the returned NIEM-M document. Search behaves exactly as if only authorized records or portions of records existed at the provider. The process of enforcing the IAP as user's access information is referred to as the *entitlement decision*.

As previously discussed in the ***Services Architecture View*** section, information providers may choose to publish to the ISI cache, allowing the ISI to implement information consumer services on their behalf, or to implement those services themselves. This distinction also applies to the enforcement of IAPs. Information providers, which do not publish to the ISI cache, must enforce information access policies. Whether enforced by the ISI and/or the information provider, user attributes that are necessary for making the entitlement decision are made available as discussed in above ***User Attributes*** section.

ENTITY ATTRIBUTES

Just as user attributes describe characteristics of users in a standardized way, *entity attributes* describe characteristics of trusted systems within the MISE. (*Entity* is SAML terminology for a *trusted system*. The usage of SAML in MISE is covered in the section ***MISE Certificate Authority***). And just as a user's information consumer system is responsible for user vetting, MISE Management Organization is responsible for trusted system vetting – i.e., confirming that the trusted system has met all criteria necessary for joining the MISE and that all entity attributes associated with the trusted system are accurate.

The MISE defines the following entity attributes that are involved in entitlement decisions, referred to as *entitlement entity attributes*:

- **Entity Id** – The unique identifier by which the entity (system) is known.
- **Law Enforcement Indicator** – Entity (system) authorized to handle law enforcement information.
- **Privacy Protected Indicator** – Entity (system) authorized to handle privacy-protected information.
- **Community Of Interest Indicator** – Entity (system) authorized to handle information shared by MISE community.

The *Indicators* associated with an entity is the full set of *Indicators* the entity is authorized to assign to users. This permits the limiting of the authority of a trusted system. For example, if an information consumer system requests information on behalf of a user and presents a *Privacy Protected Indicator* user attribute, but the information consumer system itself *does not* have a *Privacy Protected Indicator* entity attribute, the request will be denied by the Information Sharing Infrastructure.

The following entity attributes are examples of those defined for auditing purposes, referred to as *auditing entity attributes*. They are not relevant to entitlement decisions, but are recorded in audit logs as a record of the trusted system which accessed information:

- **Entity Name** – The name of the entity in a format suitable for display to a user.
- **Owner Agency Name** – The name of the organization or agency by which the entity is owned.
- **Owner Agency Country Code** –The country of the organization or agency by which the entity is owned and operated.

Additional entity attributes are defined providing technical and administrative point of contact information for the trusted system. Full details can be found in the MISE Attributes Specification.

AUDIT LOGGING

The ISI maintains audit logs that record all accesses to information from all information consumer systems to include: details regarding the request, specific information provided, criteria used to grant access, and *auditing attributes* tracking the trusted system and (when relevant) the user which requested the information.

Information consumer systems may request information on behalf of individual users, as well as on behalf of a group of users who all have an identical set of entitlement user attributes. In both, auditing entity attributes and auditing user attributes are recorded in audit logs. However, in the latter case when requesting information on behalf of a group, the information consumer system is responsible for revealing the information only to users with the precise set of auditing user attributes.

6.5.2. TRUST IMPLEMENTATION PERSPECTIVE

The ***Functional Perspective*** section focuses on MISE's *attribute-based access control* – on how user attributes, entity attributes, and information access policies are used to control which information is accessible by which users, and to audit actual accesses. This section focuses on the cryptographic methods used to establish trust and deliver attributes in a secure manner for making entitlement decisions.

The figure below illustrates key aspects of the trust implementation.

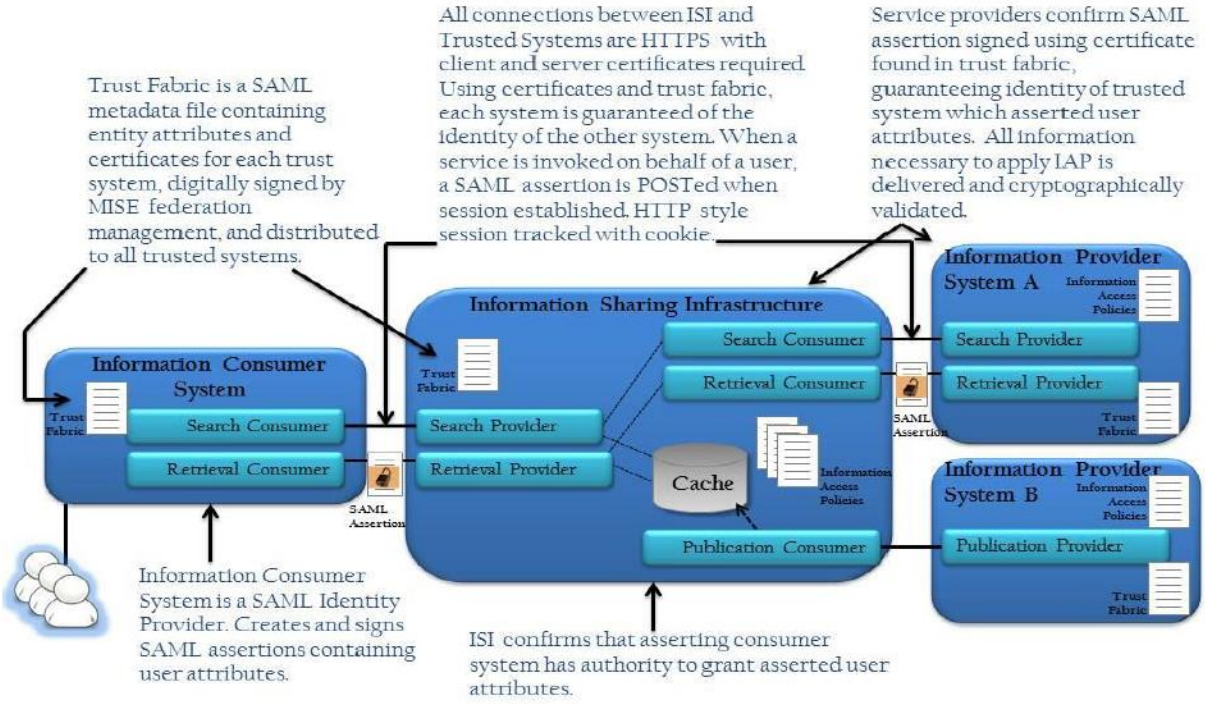

Figure 21: Security Architecture - Trust Implementation Perspective

MISE CERTIFICATE AUTHORITY

MISE Management will use a certificate authority (CA) to sign the trust fabric. Certificates signed by this certificate authority form the cryptographic foundation of trust within the MISE.

Each trusted system uses a signed certificate and matching private key to secure SSL connections. In addition, each information consumer system uses a separate signed certificate and matching private key to sign SAML assertions. The following sections provide more details on the usage of these certificates.

TRUST FABRIC

The MISE *Trust Fabric* is a digitally signed XML document in SAML metadata format describing each trusted system in the MISE. This document is signed by the MISE Certificate Authority, and distributed to all trusted systems as a trust anchor upon which all cryptographic operations rely.

In SAML metadata terms, the trust fabric document contains an *Entity Descriptor* element describing each trusted system, and also an *Entity Descriptor* element describing the ISI itself. Two types of *Role Descriptor* elements are used:

1. **Trusted System** – each *Entity Descriptor* contains a *Role Descriptor* of this type. Essential associated information includes the certificate used for SSL connections, as well as *auditing entity attributes*.
2. **Identity Provider** – information consumer systems have an identity provider role *in addition to* the trusted system role. Essential associated information includes the certificate used for signing SAML assertions, and the appropriate *Information Access Rights* entity attributes introduced above.

The organization owning each trusted system works with MISE Management to prepare the *Entity Descriptor* information describing the trusted system. When *Entity Descriptor* information is complete, has been validated for accuracy, and confirmed to meet all MISE criteria, MISE Management incorporates it into the trust fabric document, signs the document with the MISE CA private key, and distributes the updated trust fabric to all trusted systems. The same process is used when any aspect of the *Entity Descriptor* information for any trusted system changes.

The trust fabric document is made available for retrieval at any time by trusted systems from a well-known URL on the ISI, HTTPS is required to GET the trust fabric document, but a client certificate is not required. All trusted systems must retrieve the trust fabric on a periodic basis, to ensure they are always operating with the current version of the trust fabric. MISE is able to use this simple distribution mechanism for the trust fabric since the security of the MISE does not rely on the contents of the trust fabric document being kept secret, but rather upon its accuracy being guaranteed by the CA signature.

NETWORK CONNECTIONS

All network connections between the ISI and trusted systems are over the public Internet, using HTTP over SSL (HTTPS). Client and server certificates, digitally signed by the MISE certificate authority, are required for each SSL connection.

When a network connection is established, the ISI and trusted system each use the certificate submitted in the trust fabric, and makes the associated private key available to SSL software. For the purpose of SSL connections for MISE service interactions, all trusted systems must install the MISE CA cert as a trusted root CA. After the connection is established, each side can then obtain the certificate used by the other side from SSL software, and confirm it is found in the trust fabric. If it is not, the connection should be immediately dropped. Possession of the private key associated with a certificate signed by the MISE CA is thus proven, and lookup in the trust fabric proves the identity of the trusted system and that it is a current validated member of the MISE. Note that among its other roles, the trust fabric serves the purpose of the certificate revocation list, which is a part of some PKI arrangements.

This type of network connection and certificate validation allows certain security and authentication issues to be addressed at connection time, allowing simple RESTful style application level message exchanges. Specifically, the following are addressed:

- **Trusted System Authentication** – For every message exchange between the ISI and a trusted system, each system will verify the identity of the other system.
- **Message Non-repudiation and Integrity** – Every message received by the ISI or a trusted system will be verified to prove the Trusted System that claimed to have sent the message actually sent it and the message was not altered since it left control of the Trusted System that sent it, otherwise the message will be rejected.
- **Message Confidentiality** – Every message sent or received by the ISI will be cryptographically protected from being read by any person or entity other than its intended receiving system.

USER ATTRIBUTE CONVEYANCE

This section describes the mechanics of delivering authenticated user attributes to the right location at the right time, allowing entitlement decisions to be made.

Information consumer systems associate a SAML assertion with every request for information from the MISE. Information consumer systems are the authoritative source of user attributes, since they are responsible for user vetting and authentication. In SAML terms, information consumer systems act as *identity providers*. This does not preclude information consumer systems from relying on an external identity provider, including acting as a relying party to another SAML identity provider.

Characteristics of this SAML assertion include:

- Contains the appropriate authenticated set of entitlement user attributes and auditing user attributes.
- Digitally signed by the information consumer system using the private key associated with the signing certificate, which is a part of its *Identity Provider* role information within the trust fabric document. This provides a cryptographic guarantee to the ISI and to information provider systems of the identity of the information consumer system asserting the user attributes and promising to handle the information in accordance with MISE guidelines.
- Contains a SAML Audience Restriction specifying the entire MISE, meaning it is intended for interpretation by the ISI and also by information provider systems.

The figure below illustrates the mechanism used for associating this SAML assertion with a series of RESTful HTTP request / response messages making up service interactions.

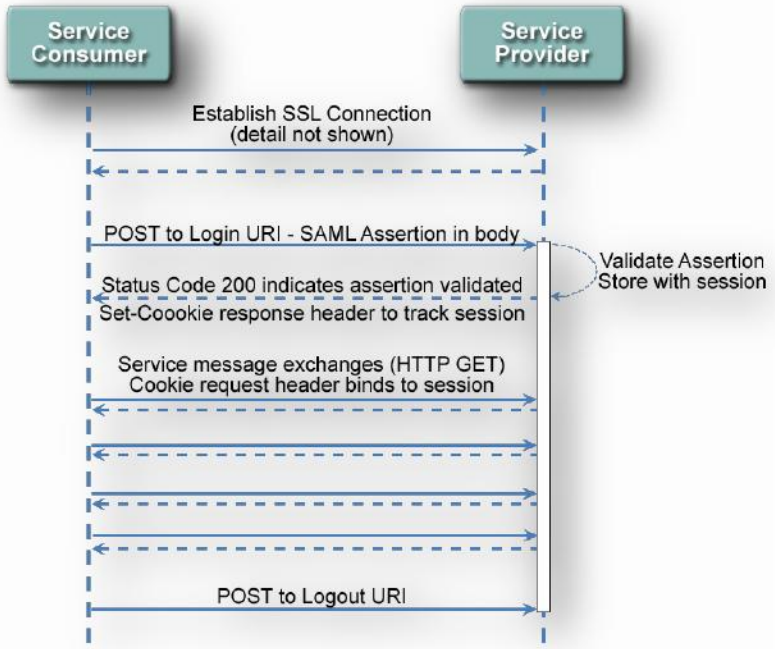

Figure 22: SAML assertion bound to HTTP session

Notice the following from the Figure 22:

- Specifics mentioned in this section occur inside an SSL connection with both systems authenticated as discussed in ***MISE Certificate Authority*** section.
- Before beginning the series of messages making up service interactions (generally a series of HTTP GETs), the service consumer does an HTTP POST to a login URI at the service provider. The body of this POST request contains the signed SAML assertion. During processing of the POST, the service provider validates the assertion. This includes standard SAML validation as defined in the MISE Interface Security Specification, as well as confirming that the signing certificate exists within the current trust fabric. If the service provider is the ISI, another validation step is ensuring that no *indicators* are assigned which the information consumer system is not authorized to assign.
- If the assertion is valid, the service provider stores the assertion in a session, returns a Set-Cookie response header to track the session, and returns an HTTP status code of 200 (OK). If the assertion is invalid, the service provider returns an HTTP status code of 403 (Forbidden).
- Once the assertion is validated and a session is created, the service consumer can proceed to interact with the provider in accordance with The MISE Search/Retrieve Interface Specification. A cookie request header specifying the session cookie assigned by the provider must be sent with each HTTP request. No additional authorization overhead is borne during each request/response. Normal web browser style session handling and expiration are used, even though this is a system-to-system interaction.

- If the service consumer knows it is finished interacting in the context of the particular set of user attributes, it can POST to a logout URI as a courtesy to the service provider, allowing the provider to delete the session and free up associated resources. If the service consumer does not explicitly logout, the service provider will delete the session automatically after a period of non-use, or if the time condition within the SAML assertion (NotOnOrAfter) is exceeded.

If the ISI is acting as a request broker, in which case the ISI acts as a service consumer to pass the request to an information provider, the same mechanism is used between the ISI and the information provider. The signed SAML assertion received from the information consumer is simply passed verbatim to the information provider, allowing the provider to identify the information consumer system which originated the request and asserted the user attributes, by locating the certificate used to sign the assertion in the trust fabric.

6.6. TECHNICAL OPERATIONS ARCHITECTURE VIEW

The Technical Operations Architecture View defines the approach for operations and sustainment of services that comprise the MISE. The primary goal of the MISE Technical Operations Architecture View is to provide a high-level overview of MISE operational concepts in the following major areas:

- Service Management
- Organizational Roles and Responsibilities
- End User Support

6.6.1. SERVICE MANAGEMENT

Service management is essential for the MISE to achieve reliable performance and meet stakeholder requirements for dissemination of critical maritime information. Service management includes activities such as monitoring availability of information provider and ISI services, registration of trusted systems for authentication and discovery services, and review of services levels to ensure service performance thresholds are met.

SERVICE LEVEL AGREEMENTS

Service level agreements (SLA) are necessary for each trusted system that connects to the ISI. The service level agreement is a formal agreement that describes the service details, documents service level targets, and specifies the responsibilities of the Trusted System and the ISI.

6.6.2. ORGANIZATIONAL ROLES AND RESPONSIBILITIES

MISE Management

The MISE Management is responsible for the day-to-day operations for the MISE. The MISE Management provides operational service management functions for the MISE and serves as the primary point of contact for the MISE Help Desk. The following list includes some specific responsibilities and functions of the MISE Management:

- Certify and approve new trusted systems to participate in the environment
- Manages MISE CA and Trust Fabric Document

- Responsible for change and configuration management of the MISE
- Manages the ISI
- Defines operational process based on best practices to be documented in the MISE Operational Policies and Procedures

Trusted Systems

The following list includes specific responsibilities of *ALL* Trusted Systems:

- Adopt entitlement attributes
- Protect PII
- Notify MISE Management of any changes to security policy or posture
- Trust MISE Root CA
- Prepare SAML metadata for inclusion in trust fabric

Responsibilities specific to Information Providers include:

- Implement entitlement attributes or delegate enforcement to MISE
- Maintain information access policy based on entitlement attributes within MISE
- Implement publication, or search/retrieve

Responsibilities specific to Information Consumers include:

- Vet end users for access to the environment
- Provide authentication credentials to end users
- Authenticate end users
- Generate SAML user assertions containing entitlement attributes
- Implement search/retrieve

6.6.3. END USER SUPPORT

The MISE is comprised of many member organizations, each with its own local help desk resources, it is necessary to leverage these local resources to the maximum extent possible. The primary guiding principle in the design of the help desk structure is that all problems should be solved as close to the user as possible, and with as little centralized effort as possible; however, more complex technical issue not contained to a single trusted system will require a central help desk entity. The following three-tier help desk structure and issue escalation plan provides guidelines for the MISE:

Tier 1: Local Help Desk Support

All issues encountered by users should be first reported to the users local help desk. This level of user assistance is provided by the users local department to resolve simple issues reported by users (e.g., network outages, firewall problems, and local desktop user interface issues) should be handled at this level without bringing any higher-level resources into play.

Tier 2: Trusted System Help Desk Support

For any issue that the local help desk cannot resolve, the trusted system provides help desk support to the user. Issues that can be solved at Tier 2 may include questions regarding permissions and access control policies for information resources provisioned via the ISI or application-specific issues on the information consumer system. The trusted system help desk should attempt to resolve the issue, and will contact the MISE help desk support staff if the issue relates to an information provider or the ISI.

Tier 3: MISE Help Desk Support

Any issue that cannot be resolved at Tier 2 should be escalated to the MISE Help Desk. Issues that can be solved at Tier 3 include repairing a corrupted trust fabric document or resolving a technical issue that arises between multiple trusted systems. Issues that are resolved at this level will be tracked in the MISE Management's issue tracking database.

THIS PAGE INTENTIONALLY LEFT BLANK